

# A Modularized Efficient Framework for Non-Markov Time Series Estimation

Gabriel Schamberg <sup>1</sup>, *Student Member, IEEE*, Demba Ba, *Member, IEEE*,  
and Todd P. Coleman <sup>2</sup>, *Senior Member, IEEE*

**Abstract**—We present a compartmentalized approach to finding the maximum *a posteriori* (MAP) estimate of a latent time series that obeys a dynamic stochastic model and is observed through noisy measurements. We specifically consider modern signal processing problems with non-Markov signal dynamics (e.g., group sparsity) and/or non-Gaussian measurement models (e.g., point process observation models used in neuroscience). Through the use of auxiliary variables in the MAP estimation problem, we show that a consensus formulation of the alternating direction method of multipliers enables iteratively computing separate estimates based on the likelihood and prior and subsequently “averaging” them in an appropriate sense using a Kalman smoother. As such, this can be applied to a broad class of problem settings and only requires modular adjustments when interchanging various aspects of the statistical model. Under broad log-concavity assumptions, we show that the separate estimation problems are convex optimization problems and that the iterative algorithm converges to the MAP estimate. As such, this framework can capture non-Markov latent time series models and non-Gaussian measurement models. We provide example applications involving 1) group-sparsity priors, within the context of electrophysiologic spectrotemporal estimation, and 2) non-Gaussian measurement models, within the context of dynamic analyses of learning with neural spiking and behavioral observations.

**Index Terms**—Bayesian, ADMM, convex optimization, sparsity, dynamics, filtering.

## I. INTRODUCTION

WE CONSIDER the problem of estimating a latent time series based on an underlying dynamic model and noisy measurements. Such a problem appears in a variety settings, including (but certainly not limited to) tracking [1], medical imaging [2], and video denoising [3]. Given the broad applica-

bility of this problem formulation, the underlying models that are used inevitably become increasingly complex.

Certain scenarios are well studied, such as the case of a linear system with Gaussian noise, where it is well known that the maximum *a-posteriori* (MAP) point estimate can be obtained using a Kalman smoother (KS) [4]. When introducing non-linearities, alternatives include the extended Kalman filter (EKF), which relies on linear approximations, as well as the unscented Kalman filter (UKF) [5] and Particle Filter (PF) [6], which use sample based techniques. While the EKF and UKF are well suited for a broad class of problems, they are not well suited for models with non-Gaussian noise. This is problematic for the increasingly popular problem of incorporating sparsity inducing models to latent signal estimation. These problems include exploiting sparsity in the underlying signal [7]–[10] in addition to exploiting sparsity in the signal dynamics [11]–[13]. While some of these methods utilize  $\ell_1$ -regularization to enforce sparsity at a local level and enable causal prediction, there is often knowledge of global structures, such as those favored by the group lasso [14], that dictate a need for batch-wise estimation. In such cases, the desired estimation problem deviates from the classical state estimation problem in that the underlying signal is no longer Markov. In such a scenario, there is no clear extension to the EKF, UKF, or PF that may be utilized to address the non-Markovicity of the underlying signal.

The broad scope of the problem in question dictates a need for a systematic approach to latent time series estimation for a variety of measurement models and system models. Furthermore, a solution framework that can compartmentalize these two models facilitates interchangeability and allows new regularization techniques to be easily incorporated to an estimation procedure.

We develop a framework using the alternating direction method of multipliers (ADMM) [15] that, under mild (i.e., log-concavity) assumptions, yields the MAP estimate for problems with non-Markov latent variables and/or nonlinear observations. While ADMM has been utilized to decompose specific dynamic systems into simpler subproblems [13], [16], our approach applies to arbitrary log-concave dynamic models. In particular, we utilize auxiliary variables to enable a solution involving iterative updates to three modules, one that pertains to the measurement model, another that pertains to the prior distribution on the latent signal, and a third that is a Kalman smoother. As such, our framework enables various sparsity models to be easily applied to the signal and/or dynamics with adjustments only required to the corresponding module. We demonstrate implementation of

Manuscript received June 9, 2017; revised November 7, 2017; accepted December 27, 2017. Date of publication January 15, 2018; date of current version May 2, 2018. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Wenwu Wang. The work of G. Schamberg was supported by the Jacobs Fellowship from the Department of Electrical and Computer Engineering, University of California, San Diego. The work of T. P. Coleman was supported by Grants NIH 1 R01MH110514 and NSF CCF-0939370. (*Corresponding author: Todd P. Coleman.*)

G. Schamberg is with the Department of Electrical and Computer Engineering, University of California, San Diego, CA 92092 USA (e-mail: gschambe@eng.ucsd.edu).

D. Ba is with the School Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138 USA (e-mail: demba@seas.harvard.edu).

T. P. Coleman is with the Department of Bioengineering, University of California, San Diego, CA 92092 USA (e-mail: tpcoleman@ucsd.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2018.2793870

the framework in two distinct applications, namely latent state estimation and spectrotemporal estimation. We show that in the case of state estimation, our method outperforms a fixed interval smoother and particle filter for two state-space models coupled with non-Gaussian observations. In the case of spectrotemporal estimation, we demonstrate the efficacy of our method when using non-Markov priors. The proposed method yields an intuitive approach to latent process estimation with iterative use of a Kalman smoother in tandem with standard convex optimization techniques. We provide a mathematical justification for the intuition by proving that our approach guarantees convergence to the MAP solution under the same relatively mild conditions that apply to general ADMM approaches. Finally, we provide software to enable the reader to reproduce the results of this paper and to easily apply the framework to novel models.<sup>1</sup> Our contributions may be summarized as follows:

- We present an efficient iterative solution framework for latent time series estimation with a guarantee of convergence to the MAP estimate under mild log-concavity assumptions.
- In the presence of non-Linear, non-Gaussian measurement models, our method does not require a Gaussian approximation, unlike KS variants, and is more efficient than Sequential Monte Carlo (SMC) methods.
- Our framework accommodates non-Markov signals despite there being no clear method for adapting EKF, UKF, and SMC methods for such a scenario, particularly when the prior applies to highly non-linear functions of the latent process, such as a singular value decomposition.
- Through the use of auxiliary variables, the ADMM solution to our reformulated MAP estimation problem is modular, with the observation and system models in disjoint modules that are unified by a Kalman smoother.

The paper is structured as follows: Section II provides the general formulation of the problem we are solving in addition to a brief review of relevant work solving specific instances of the problem. Section III details a novel systematic approach for solving the MAP estimation problem in its general form. Section IV demonstrates the capabilities of the framework through implementation on two existing problems. Section V concludes the paper with a discussion of the results and future work.

## II. PROBLEM FORMULATION

### A. Notation

While it is intended that the notation is presented unambiguously, we here present some notational conventions. Bold letters are used to represent vectors and matrices, whereas non-bold letters represent scalars. Subscripts are used for indexing scalar elements of a vector, or columns of a matrix. A double subscript is used to specify scalar elements of a matrix. For example,  $x_n$

<sup>1</sup>An implementation of the proposed framework can be found in the following publicly available GitHub repository: <https://github.com/gabeschamberg/nonmarkov-timeseries-estimation/releases/tag/v1.1>. This repository includes the iPython [17] notebooks that were used to generate Figs. 2 and 3 and Table III.

gives the  $n$ th element of a vector  $\mathbf{x}$ ,  $\mathbf{x}_n$  gives the  $n$ th column of a matrix  $\mathbf{x}$ , and  $x_{n,m}$  gives the  $m$ th row of the  $n$ th column of a matrix  $\mathbf{x}$ . Capital/lowercase letter pairs represent either random variable/realization pairs or total count/index pairs. For example, we may have that  $\mathbf{x}_n$  gives a specific value of the random vector  $\mathbf{X}_n$ , which is the  $n$ th column of a random matrix  $\mathbf{X}$  with  $N$  columns in total. We let  $f$  and  $p$  denote probability density functions (pdfs) and probability mass functions (pmfs), respectively. Various joint and conditional pdfs and pmfs are made clear by their subscripts. For example, the pdf of  $X$  given  $Y = y$  is  $f_{X|Y}(\cdot|y)$ . We let  $\mathbb{R}$  denote the space of real numbers,  $\mathbb{R}_+$  denote the non-negative reals,  $\mathbb{R}^{A \times B}$  denote the space of  $A$  by  $B$  real valued matrices, and  $\mathbb{R}^{AB}$  denote the space of real valued vectors of length  $A$  times  $B$ .

### B. Problem Setup

Let  $\mathcal{X}$  and  $\mathcal{Y}$  be measurable spaces and  $N$  be the length of time series pertaining to the latent process  $\mathbf{X} \in \mathcal{X}^N$  and observed process  $\mathbf{Y} \in \mathcal{Y}^N$ . Unless otherwise specified, we assume  $\mathcal{X} = \mathbb{R}^K$  and  $\mathcal{Y} = \mathbb{R}^P$  where  $K$  is the dimension of the latent process at any time, and  $P$  is the dimension of the observation process at any time. As such,  $\mathbf{X} \in \mathbb{R}^{K \times N}$  is the latent time series we wish to estimate and  $\mathbf{Y} \in \mathbb{R}^{P \times N}$  is the collection of noisy observations. Furthermore, assume that these observations are conditionally independent given the underlying time series:

$$f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y} | \mathbf{x}) = \prod_{n=1}^N f_{Y_n|\mathbf{X}_n}(y_n | \mathbf{x}_n) \quad (1)$$

where  $f_{\mathbf{Y}|\mathbf{X}}$  is the likelihood of the entire collection of observations given the entire latent time series and  $f_{Y_n|\mathbf{X}_n}$  is the likelihood of a single observation given the corresponding element of the latent time series.

Next, define the latent signal's dynamics (or system behavior) in terms of  $\mathbf{W} \in \mathbb{R}^{K \times N}$  for which

$$\mathbf{W}_n = \begin{cases} \mathbf{X}_1 & n=1 \\ \mathbf{X}_n - \mathbf{D}\mathbf{X}_{n-1} & n=2, \dots, N \end{cases},$$

where  $\mathbf{D} \in \mathbb{R}^{K \times K}$  is a transition matrix and  $\mathbf{W}_n \in \mathbb{R}^K$  and  $\mathbf{X}_n \in \mathbb{R}^K$  represent the  $n$ th columns of  $\mathbf{W}$  and  $\mathbf{X}$ , respectively. For compactness we write this as  $\mathbf{W} = \mathcal{A}(\mathbf{X})$ , where  $\mathcal{A}$  represents a linear operator that is fully defined by  $\mathbf{D}$ . We assume that  $\mathbf{W}$  is distributed according to a known prior pdf  $f_{\mathbf{W}}(\mathbf{w})$ . Note that this framework includes, for the special case of  $\mathbf{W}_n = \mathbf{X}_n - \mathbf{X}_{n-1}$  and  $\mathbf{W}_n \sim \mathcal{N}(\mu_n, \Sigma_n)$  are independent Gaussian random vectors for  $n = 2, \dots, N$ , the well-studied scenario in which the underlying time series  $\mathbf{X}$  is a Gauss-Markov process.

Here, we consider the problem of finding the maximum a posteriori estimate:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} -\log f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y} | \mathbf{x}) - \log f_{\mathbf{X}}(\mathbf{x}) \quad (2)$$

where  $-\log f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y} | \mathbf{x})$  is the negative log-likelihood and  $-\log f_{\mathbf{X}}(\mathbf{x})$  is the negative log-prior. We note that because  $\mathbf{W}$  is a linear function of  $\mathbf{X}$ , we have  $f_{\mathbf{X}}(\mathbf{x}) \propto f_{\mathbf{W}}(\mathcal{A}(\mathbf{x}))$ . This relationship indicates that knowing a prior on either  $\mathbf{X}$  or  $\mathbf{W}$

induces a prior on the other. Thus, we can equivalently rewrite our problem as:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} -\log f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y} | \mathbf{x}) - \log f_{\mathbf{W}}(\mathcal{A}(\mathbf{x})) \quad (3)$$

$$= \underset{\mathbf{x}}{\operatorname{argmin}} L(\mathbf{y} | \mathbf{x}) + \beta\phi(\mathcal{A}(\mathbf{x})) \quad (4)$$

with  $\beta \in \mathbb{R}_+$  and where we define the measurement model  $L : \mathbb{R}^{N \times K} \rightarrow \mathbb{R}$  and system model  $\phi : \mathbb{R}^{N \times K} \rightarrow \mathbb{R}$  as:

$$L(\mathbf{y} | \mathbf{x}) := -\log f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y} | \mathbf{x}) \quad (5)$$

$$\phi(\mathbf{w}) := -\frac{\log f_{\mathbf{W}}(\mathbf{w})}{\beta}. \quad (6)$$

The inclusion of  $\beta$  in (4) is to facilitate the cases when the system model is only known up to a proportionality constant or when  $\phi$  is a regularizer used to exploit a desired dynamic characteristic of the latent signal (as opposed to representing the *true* distribution of  $\mathbf{W}$ ). In either of these cases  $\beta$  is interpreted as a tuning parameter used to control the extent to which the system model is weighted (as in  $\lambda$  throughout [14]).

Throughout this paper, we will interchangeably use the names log-likelihood/measurement model in reference to  $L$ , and log-prior/system model/dynamic model in reference to  $\phi$ . Due to the assumption that observations are conditionally independent given the state variables, the measurement model can be decomposed into a sum over  $N$  measurements, each depending on the state variable at a single time instance:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \left( \sum_{n=1}^N L_n(\mathbf{y}_n | \mathbf{x}_n) \right) + \beta\phi(\mathcal{A}(\mathbf{x})) \quad (7)$$

where  $L_n(\mathbf{y}_n | \mathbf{x}_n) := -\log f_{\mathbf{Y}_n|\mathbf{X}_n}(\mathbf{y}_n | \mathbf{x}_n)$ . It should be noted that the problem presented in (7) is made difficult by the second term. In particular, imposing a prior on the *differences* of the underlying time series prevents separability across the  $N$  time points. Furthermore, by allowing for non-Markov models, it is possible to have models that do not allow the second term to be separated into terms each containing only  $\mathbf{x}_n$  and  $\mathbf{x}_{n-1}$  for each  $n = 1 \dots N$ . In the following section, we present a framework for efficiently solving problems in the form of (7) for a broad class of measurement models  $L$  and system models  $\phi$ .

### C. Related Work

Works related to our proposed method include both the investigation of new algorithms for estimating latent time series and the creation/application of new time series models. Notably, the Kalman smoother [4] and its variants [5], [6] provide structured approaches to estimating latent signals in a subset of problems with dynamical system models and noisy measurements. While the Kalman smoother is MAP optimal for the very specific case of a linear system with Gaussian noise, its non-linear variants do not guarantee optimality and do not offer solutions for a comprehensive class of measurement and system models. In particular, there has been growing interest in models exploiting the sparsity of states and/or dynamics of signals [7]–[9], [11]–[13], which in many cases do not lend themselves to solutions via the existing Kalman smoother variants.

TABLE I

EXAMPLES OF COMMON MODELS. FOR THE MULTIPLE MODALITIES CASE, WE DEFINE  $\mathbf{y} = (\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(J)})$  TO BE A  $J$ -TUPLE OF SIMULTANEOUS AND CONDITIONALLY INDEPENDENT OBSERVATIONS, EACH WITH ITS OWN DIMENSIONALITY AND ASSOCIATED MEASUREMENT MODEL  $L^{(j)}$

Measurement Models - $L(\mathbf{y}   \mathbf{x})$	
Linear Gaussian (LG)	$\sum_{n=1}^N \ \mathbf{y}_n - \mathbf{A}\mathbf{x}_n + \mathbf{b}\ _2^2 \triangleq LG(\mathbf{x})$
Sparse LG	$LG(\mathbf{x}) + \ \mathbf{x}\ _1$
Group sparse LG	$LG(\mathbf{x}) + \sum_{k=1}^K \left( \sum_{n=1}^N x_{k,n}^2 \right)^{\frac{1}{2}}$
Multiple Modalities	$\sum_{j=1}^J L^{(j)}(\mathbf{y}^{(j)}   \mathbf{x})$
System Models - $\phi(\mathbf{w})$	
LG	$\sum_{n=1}^N \ \mathbf{C}\mathbf{w}_n - \mathbf{d}\ _2^2$
Sparse	$\ \mathbf{w}\ _1$
Group sparse	$\sum_{k=1}^K \left( \sum_{n=1}^N w_{k,n}^2 \right)^{\frac{1}{2}}$

TABLE II

EXAMPLES OF MEASUREMENT MODEL/SYSTEM MODEL PAIRINGS IN PREVIOUS WORKS

	Measurement Model	System Model
<i>Kalman Smoother</i> [4]	LG	LG
<i>State Space Model of Learning</i> [21]	Non-linear/ multiple modalities	Gaussian
<i>Spectrotemporal Pursuit</i> [11]	LG	Group sparse
<i>Lasso-Kalman Smoother</i> [13]	Group sparse	LG
<i>Sparse States and Sparse Innovations</i> [12]	Sparse	Sparse

For such sparsity-inducing models, existing causal estimators are often heuristic extensions of the Kalman filter, such as  $\ell_1$ -regularized Kalman filter updates [12] and tracking a belief of the support set [7]. Causal estimation is made particularly challenging for the models that are non-Markov in nature. As such, the aforementioned causal estimators lack performance guarantees. Existing batchwise solutions utilize a Kalman smoother to solve the updates for a particular iterative algorithm, such as IRLS for group sparse dynamics [11] and ADMM for group sparse states [13]. In the latter example, their non-consensus formulation of ADMM is reliant upon the choice of a Gaussian system model.

In addition to the Kalman smoother variants, sample based methods such as Markov chain Monte Carlo (MCMC) and SMC are viable options for latent time series estimation. While these methods can accommodate non-linear and non-Gaussian models [18] and can simultaneously estimate the state and model parameters [19], [20], they are often computationally prohibitive. Furthermore, these methods do not have a straightforward extension to non-Markov and non-linear priors such as the  $\ell_1/\ell_2$  and nuclear norm priors (see Remark 2).

Here we propose a generalized framework for obtaining the MAP estimate in many of the aforementioned problems in a batchwise manner. Tables I and II show the models used in some of these problems and serve to illustrate the primary contribution of our framework, namely that for a given problem, the solution is modular in that the choice of measurement model can be

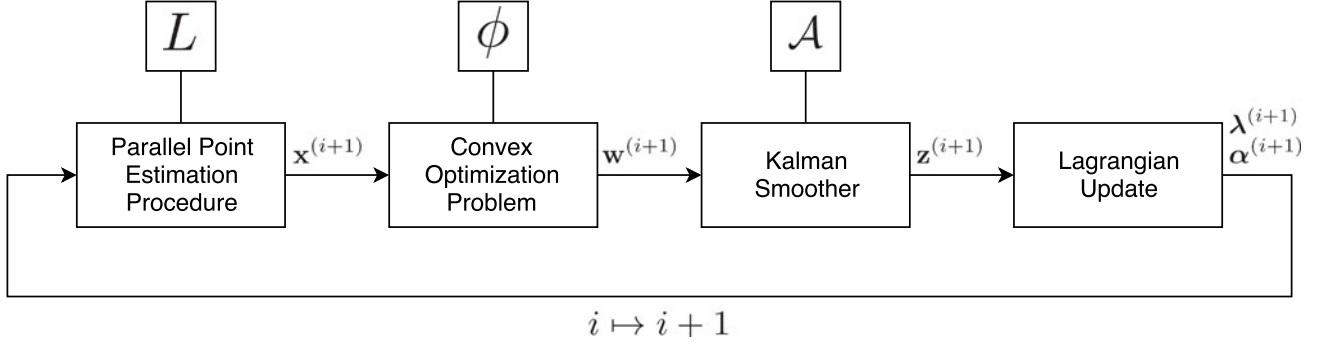


Fig. 1. Block diagram of the modular MAP estimation framework illustrates how the selection of  $L$ ,  $\phi$ , and  $\mathcal{A}$  affects independent parts of the estimation procedure.

made independently of the system model without requiring a complete rederivation of the solution.

### III. MODULAR MAP ESTIMATION FRAMEWORK

The alternating direction method of multipliers (ADMM) allows large global problems to be decomposed into smaller sub-problems whose solutions can be coordinated to achieve the global solution. ADMM offers an iterative solution of the dual problem that has the decomposability of dual descent in addition to the convergence guarantees of the method of multipliers, which hold under fairly mild conditions. While the details of dual optimization and ADMM are omitted here, they can be found in [22] and [15], respectively.

We begin by reformulating (7) to create separability in the objective function by including  $\mathbf{w}$  as an optimization variable and introducing a constraint to preserve the relationship between  $\mathbf{x}$  and  $\mathbf{w}$ :

$$\begin{aligned} (\hat{\mathbf{x}}, \hat{\mathbf{w}}) &= \underset{\mathbf{x}, \mathbf{w}}{\operatorname{argmin}} \sum_{n=1}^N L_n(\mathbf{y}_n | \mathbf{x}_n) + \beta \phi(\mathbf{w}) \\ \text{s.t.} \quad & \mathbf{w} = \mathcal{A}(\mathbf{x}). \end{aligned} \quad (8)$$

The optimization problem given by (8) can be solved using ADMM, and would yield a solution that enables the measurement model and penalty function to be addressed in independent subproblems. However, when using the above formulation, the update equations yielded by the ADMM algorithm would require one of the aforementioned approximate or sample-based methods for non-Gaussian measurement models (see Remark 1).

We use a variant of ADMM known as consensus ADMM and construct a modular solution framework shown in Fig. 1 that only requires making local adjustments to the solution when modifying the measurement model ( $L$ ), penalty function ( $\phi$ ), or transition model ( $\mathcal{A}$ ). This is accomplished by introducing an auxiliary variable  $\mathbf{z} \in \mathbb{R}^{K \times N}$  to achieve separability (of  $\mathbf{x}$  and  $\mathbf{w}$ ) in the constraints as well as the objective function:

$$\begin{aligned} (\hat{\mathbf{x}}, \hat{\mathbf{w}}, \hat{\mathbf{z}}) &= \underset{\mathbf{x}, \mathbf{w}, \mathbf{z}}{\operatorname{argmin}} \sum_{n=1}^N L_n(\mathbf{y}_n | \mathbf{x}_n) + \beta \phi(\mathbf{w}) \\ \text{s.t.} \quad & \mathbf{x} = \mathbf{z} \\ & \mathbf{w} = \mathcal{A}(\mathbf{z}). \end{aligned} \quad (9)$$

The optimization problem given by (9) is termed the consensus formulation, and  $\mathbf{z}$  the consensus variable. By introducing this variable, our iterative updates with respect to the measurement model and penalty function are not only independent of each other, but are also independent of the transition model determined by  $\mathcal{A}$ .

The first step in solving (9) using ADMM requires generating the augmented Lagrangian:

$$\begin{aligned} \mathcal{L}_\rho(\mathbf{x}, \mathbf{w}, \mathbf{z}, \lambda, \alpha) &= \sum_{n=1}^N L_n(\mathbf{y}_n | \mathbf{x}_n) + \beta \phi(\mathbf{w}) \\ &+ \langle \lambda, \mathbf{x} - \mathbf{z} \rangle + \langle \alpha, \mathbf{w} - \mathcal{A}(\mathbf{z}) \rangle \\ &+ \frac{\rho}{2} \|\mathbf{x} - \mathbf{z}\|_F^2 + \frac{\rho}{2} \|\mathbf{w} - \mathcal{A}(\mathbf{z})\|_F^2 \end{aligned} \quad (10)$$

where  $\lambda \in \mathbb{R}^{K \times N}$  and  $\alpha \in \mathbb{R}^{K \times N}$  are Lagrange multipliers,  $\langle \cdot, \cdot \rangle$  is the Frobenius inner product,  $\|\cdot\|_F$  is the matrix Frobenius norm, and  $\rho \in \mathbb{R}_+$  is the penalty parameter for the augmented Lagrangian. Note that in the case where  $\rho = 0$ , the augmented Lagrangian is equivalent to the standard (unaugmented) Lagrangian.

Given the augmented Lagrangian, the ADMM solution is obtained by iteratively alternating between minimization with respect to the primal variables ( $\mathbf{x}$ ,  $\mathbf{w}$  and  $\mathbf{z}$ ) and performing gradient ascent on the Lagrange multipliers. These iterations represent a trade off between finding a solution that minimizes the cost function in (9) while ensuring that the Lagrange multipliers are such that the dual function of (9) is increasing in  $i$  and thus ensuring the constraints are satisfied. Letting  $\mathbf{x}^{(i)}$  represent the estimate of  $\mathbf{x}$  after  $i$  iterations (similarly for  $\mathbf{w}^{(i)}$ ,  $\mathbf{z}^{(i)}$ ,  $\lambda^{(i)}$ , and  $\alpha^{(i)}$ ), each iteration of ADMM is composed of the following updates [15, Sec. 3.1]:

$$\begin{aligned} \mathbf{x}^{(i+1)} &= \underset{\mathbf{x}}{\operatorname{argmin}} \mathcal{L}_\rho(\mathbf{x}, \mathbf{w}^{(i)}, \mathbf{z}^{(i)}, \lambda^{(i)}, \alpha^{(i)}) \\ \mathbf{w}^{(i+1)} &= \underset{\mathbf{w}}{\operatorname{argmin}} \mathcal{L}_\rho(\mathbf{x}^{(i+1)}, \mathbf{w}, \mathbf{z}^{(i)}, \lambda^{(i)}, \alpha^{(i)}) \\ \mathbf{z}^{(i+1)} &= \underset{\mathbf{z}}{\operatorname{argmin}} \mathcal{L}_\rho(\mathbf{x}^{(i+1)}, \mathbf{w}^{(i+1)}, \mathbf{z}, \lambda^{(i)}, \alpha^{(i)}) \\ \lambda^{(i+1)} &= \lambda^{(i)} + \rho(\mathbf{x}^{(i+1)} - \mathbf{z}^{(i+1)}) \\ \alpha^{(i+1)} &= \alpha^{(i)} + \rho(\mathbf{w}^{(i+1)} - \mathcal{A}(\mathbf{z}^{(i+1)})). \end{aligned} \quad (11)$$

By fixing all but one variable in each update, the objective functions can be simplified by dropping the terms in (10) that do not contain the optimization variable for the corresponding update. As a result, when updating with respect to the measurement model  $L$  and the system model  $\phi$ , we only need to consider the model corresponding to that update and an  $\ell_2$ -norm proximal operator [23] that ensures the update is moving in the appropriate direction to achieve a global consensus. This inclusion of the proximal operators in the augmented Lagrangian enables the use of ADMM with non-smooth objective functions [23, Sec 4.4]. Then, updating of the consensus variable involves “centering” it such that it gives equal representation to our current estimates based on the measurements and our estimates based on the system dynamics. In this sense, our ADMM framework yields a mathematical justification for a very intuitive approach, namely, iteratively finding the best estimate based on measurements, finding the best estimate based on dynamics, and “averaging” the two in the appropriate sense. This viewpoint will be made clearer in the following sections where we detail the specific update equations.

#### A. Measurement Model Update

When updating with respect to the measurement model, only terms containing  $\mathbf{x}$  in the augmented Lagrangian must be considered. To simplify notation, we will consider the scaled form of the update equations [15, Sec. 3.1.1], which can be obtained by combining the appropriate linear and quadratic terms in (10) by completing the square:

$$\mathbf{x}^{(i+1)} = \underset{\mathbf{x}}{\operatorname{argmin}} \sum_{n=1}^N L_n(\mathbf{y}_n | \mathbf{x}_n) + \frac{\rho}{2} \|\mathbf{x} - \tilde{\mathbf{x}}^{(i)}\|_F^2 \quad (12)$$

where  $\tilde{\mathbf{x}}^{(i)} := \mathbf{z}^{(i)} - \lambda^{(i)}/\rho$  is fixed within the scope of this update. Details for deriving the scaled form of the update can be found in Appendix A. Given that the squared Frobenius norm can be decomposed to the sum of squared  $\ell_2$  norms, we note that the measurement model update is separable over  $n$ , meaning that we can solve for  $\mathbf{x}_n^{(i+1)}$  for each  $n = 1, \dots, N$  independently:

$$\mathbf{x}_n^{(i+1)} = \underset{\mathbf{x}_n}{\operatorname{argmin}} L_n(\mathbf{y}_n | \mathbf{x}_n) + \frac{\rho}{2} \|\mathbf{x}_n - \tilde{\mathbf{x}}_n^{(i)}\|_2^2 \quad (13)$$

where  $\tilde{\mathbf{x}}_n^{(i)} := \mathbf{z}_n^{(i)} - \lambda_n^{(i)}/\rho$ .

*Remark 1:* Note that the ability to separate each of the  $N$  updates is a result of the inclusion of the consensus variable. Excluding this variable would require that the dynamics be considered in the update of the measurement model:

$$\begin{aligned} \mathbf{x}^{(i+1)} = \underset{\mathbf{x}}{\operatorname{argmin}} & \sum_{n=1}^N L_n(\mathbf{y}_n | \mathbf{x}_n) \\ & + \frac{\rho}{2} \|\mathbf{x}_n - \mathbf{D}\mathbf{x}_{n-1} - \tilde{\mathbf{x}}_n^{(i)}\|_2^2 \end{aligned}$$

where  $\tilde{\mathbf{x}}_n^{(i)} := \mathbf{w}_n^{(i)} - \gamma_n^{(i)}/\rho$ ,  $\mathbf{x}_0 := 0$ , and  $\gamma$  represents the single Lagrange multiplier that would be required in solving (8) using ADMM. Requiring that the dynamics of the underlying time series be included in the measurement model update prohibits

solving for  $\mathbf{x}_n^{(i)}$  independently across  $N$ . Thus, using ADMM in this fashion does not offer any simplifications over traditional approaches for non-Gaussian measurement models. As such, incorporation of the consensus variable not only enables faster processing by allowing each update to be parallelized across  $N$ , but it allows the framework to be applied in a straightforward, non-approximate manner to a broad class of measurement models.

It should be noted that while we assume conditional independence of the observations given the latent time series, one can revert to the update in (12) for the case when the observations are correlated. In this case the ability to parallelize across  $n$  is lost, but the ability to ignore system dynamics is preserved (i.e., the optimization problem in (12) still does not depend on  $\phi$ ).

#### B. System Model Update

In the system model update, only terms in (10) that contain  $\mathbf{w}$  must be included. Again, we consider the scaled form:

$$\mathbf{w}^{(i+1)} = \underset{\mathbf{w}}{\operatorname{argmin}} \beta\phi(\mathbf{w}) + \frac{\rho}{2} \|\mathbf{w} - \tilde{\mathbf{w}}^{(i)}\|_F^2 \quad (14)$$

where  $\tilde{\mathbf{w}}^{(i)} := \mathcal{A}(\mathbf{z}^{(i)}) - \alpha^{(i)}/\rho$ . In this form we can clearly interpret the system model update as finding a new collection of latent variable transitions  $\mathbf{w}^{(i+1)}$  that is both representative of our system model  $\phi$  and proximal to the appropriately scaled current consensus on the transitions  $\tilde{\mathbf{w}}^{(i)}$ .

The key observation is that this framework imposes no restrictions as to whether or not our underlying signal is Markov. In the case where the signal is indeed Markov, then  $\mathbf{w}_n^{(i+1)}$  would be updated independently over  $n$ , but in general we do not assume this is the case. This provides the ability to impose batch-level structures on the dynamics of the signal. Furthermore, we note that the nature of the proximal operator enables closed form solutions when  $\phi$  is chosen to be a number of common sparsity inducing priors. In particular, because the proximal operator is not multiplying  $\mathbf{w}$  by a non-orthonormal matrix, the  $\ell_1$ , group sparse, and nuclear norm priors all offer soft-thresholding solutions [24]. Furthermore, we note that for a fixed  $K$ , the complexity of the soft-thresholding solutions for the  $\ell_1$  and group sparse priors scale linearly with  $N$  per iteration. The nuclear norm prior, however, requires a singular value decomposition (SVD), and thus scales quadratically with  $N$  per iteration [25]. Similarly, for a fixed  $N$ , the same scaling factors apply to  $K$ . It should be noted however, that if increasing  $N$  and  $K$ , the complexity of the SVD will scale quadratically with  $\max\{K, N\}$  and cubically with  $\min\{K, N\}$ .

#### C. Consensus Update

Updating the consensus variable depends on neither the measurement model nor the system model. We can think of this step as averaging our current estimates of our signal based on measurements  $\mathbf{x}^{(i+1)}$  and based on dynamics  $\mathbf{w}^{(i+1)}$ :

$$\mathbf{z}^{(i+1)} = \underset{\mathbf{z}}{\operatorname{argmin}} \|\mathbf{z} - \tilde{\mathbf{z}}_x^{(i)}\|_F^2 + \|\mathcal{A}(\mathbf{z}) - \tilde{\mathbf{z}}_w^{(i)}\|_F^2 \quad (15)$$

where  $\tilde{\mathbf{z}}_{\mathbf{x}}^{(i)} := \mathbf{x}^{(i+1)} + \lambda^{(i)}/\rho$  and  $\tilde{\mathbf{z}}_{\mathbf{w}}^{(i)} = \mathbf{w}^{(i+1)} + \boldsymbol{\alpha}^{(i)}/\rho$ . Note that given the nature of the linear operator  $\mathcal{A}$ , (15) can always be solved efficiently using a Kalman smoother.

This step clarifies the notion of ‘‘averaging’’ the current estimates  $\mathbf{x}^{(i+1)}$  and  $\mathbf{w}^{(i+1)}$ . By framing our problem from a consensus ADMM perspective, we can carve out various elements of the model and delegate them to independent updates. Then, given the nature of the relationship between the signal  $\mathbf{x}$  and the dynamics  $\mathbf{w}$ , establishing consensus between the two estimates is a Kalman smoothing problem *regardless* of the measurement and system models. This is a result of the use of  $\ell_2$ -norms in the augmented Lagrangian, which can be thought of as representing Gaussian noise with identity covariance. In other words, at each iteration  $i$ , the consensus update is a Kalman smoothing problem where each of our measurements are given by  $\tilde{\mathbf{z}}_{\mathbf{x}}^{(i)}$  and each of our predictions are given by  $\tilde{\mathbf{z}}_{\mathbf{w}}^{(i)}$ . In this sense, the consensus update gives equal weight to the current iterates of our measurement and system estimates. This follows from the fact that the log-likelihood and log-prior have their own uncertainty terms that dictate how far the updates  $\mathbf{x}^{(i+1)}$  and  $\mathbf{w}^{(i+1)}$  can deviate from the consensus in their respective updates, namely measurement noise and the tuning parameter  $\beta$ . We note that because both terms in (15) can be thought of as representing Gaussian noise with identity covariance and the transition model  $\mathcal{A}$  is invariant over iterations  $i$ , all matrix inversions required by the Kalman smoother can be precomputed. As a result, each iteration requires on the order of  $N$  matrix multiplications.

#### D. Convergence

Next we consider the practical and theoretical convergence of the proposed framework. To begin, we present the optimality conditions and the means with which we can in practice implement convergence checks. The derivations are omitted, as they closely follow Section 3.3 of [15]. The optimality conditions for the proposed framework are given by:

$$\left. \begin{aligned} 0 &= \hat{\mathbf{x}} - \hat{\mathbf{z}} \\ 0 &= \hat{\mathbf{w}} - \mathcal{A}(\hat{\mathbf{z}}) \end{aligned} \right\} \text{Primal Feasibility}$$

$$\left. \begin{aligned} 0 &\in \frac{\partial}{\partial \hat{\mathbf{x}}} L(\mathbf{y} | \hat{\mathbf{x}}) + \hat{\lambda} \\ 0 &\in \frac{\partial}{\partial \hat{\mathbf{w}}} \beta \phi(\hat{\mathbf{w}}) + \hat{\boldsymbol{\alpha}} \\ 0 &= \hat{\lambda} + \tilde{\mathcal{A}}(\hat{\boldsymbol{\alpha}}) \end{aligned} \right\} \text{Dual Feasibility} \quad (16)$$

where  $\partial/\partial \cdot$  is the subgradient operator (or gradient when defined, in which case  $\in$  becomes an equality) and where  $\tilde{\mathcal{A}}(\mathbf{A})_n = \mathbf{A}_n - \mathbf{D}^T \mathbf{A}_{n+1}$  for  $n = 1, \dots, N-1$  and  $\tilde{\mathcal{A}}(\mathbf{A})_N = \mathbf{A}_N$  for  $A \in \mathbb{R}^{K \times N}$ . The primal feasibility conditions ensure that our  $\hat{\mathbf{z}}$  preserves the desired relationship between  $\hat{\mathbf{x}}$  and  $\hat{\mathbf{w}}$ , and the dual feasibility conditions serve the purpose of ensuring that the optimal Lagrange multipliers are such that  $\hat{\mathbf{x}}$  and  $\hat{\mathbf{w}}$  jointly minimize  $L$  and  $\phi$ .

Using these optimality conditions, we can derive the primal and dual residuals:

$$\left. \begin{aligned} r_1^{(i)} &= \mathbf{x}^{(i)} - \mathbf{z}^{(i)} \\ r_2^{(i)} &= \mathbf{w}^{(i)} - \mathcal{A}(\mathbf{z}^{(i)}) \end{aligned} \right\} \text{Primal Residuals}$$

$$\left. \begin{aligned} s_1^{(i)} &= \rho \tilde{\mathcal{A}}(\mathbf{w}^{(i)} - \mathbf{w}^{(i-1)}) \\ s_2^{(i)} &= \rho(\mathbf{z}^{(i)} - \mathbf{z}^{(i-1)}) \end{aligned} \right\} \text{Dual Residuals} \quad (17)$$

where primal feasibility is achieved when  $r_j^{(i)} = 0$  and dual feasibility is achieved when  $s_j^{(i)} = 0$  for all  $j \in \{1, 2\}$ . In practice, we declare the algorithm converged when  $\|r_j^{(i)}\|_F \leq \epsilon_j^{pri}$  and  $\|s_j^{(i)}\|_F \leq \epsilon_j^{dual}$  for all  $j \in \{1, 2\}$ , with the thresholds given by:

$$\begin{aligned} \epsilon_1^{pri} &= \epsilon^{rel} \max\{\|\mathbf{x}^{(i)}\|_F, \|\mathbf{z}^{(i)}\|_F\} + \epsilon^{abs} \sqrt{KN} \\ \epsilon_2^{pri} &= \epsilon^{rel} \max\{\|\mathbf{w}^{(i)}\|_F, \|\mathcal{A}(\mathbf{z}^{(i)})\|_F\} + \epsilon^{abs} \sqrt{KN} \\ \epsilon_1^{dual} &= \epsilon^{rel} \|\lambda^{(i)}\|_F + \epsilon^{abs} \sqrt{KN} \\ \epsilon_2^{dual} &= \epsilon^{rel} \|\boldsymbol{\alpha}^{(i)}\|_F + \epsilon^{abs} \sqrt{KN} \end{aligned} \quad (18)$$

where  $\epsilon^{rel}$  (relative tolerance) and  $\epsilon^{abs}$  (absolute tolerance) are small positive parameters.

In general, ADMM does not guarantee convergence for more than two optimization variables [26]. As such, it is not immediately clear that our ADMM framework would guarantee convergence given that it optimizes over  $\mathbf{x}$ ,  $\mathbf{w}$ , and  $\mathbf{z}$ . As it turns out, for the particular version of consensus ADMM that we are proposing, we can guarantee convergence under the same mild conditions required in standard ADMM.

*Theorem 1:* Given an observation  $\mathbf{y}$ , when  $L(\mathbf{y} | \cdot)$  and  $\phi(\cdot)$  are closed, proper, and convex functions, the ADMM algorithm given by (10) and (11) converges to the solution of (9), i.e.,  $(\mathbf{x}^{(i)}, \mathbf{w}^{(i)}, \mathbf{z}^{(i)}) \rightarrow (\hat{\mathbf{x}}, \hat{\mathbf{w}}, \hat{\mathbf{z}})$  as  $i \rightarrow \infty$ .

The proof of Theorem 1 is based on a consensus ADMM formulation presented in section 5 of [27] and is given in detail in Appendix B.

## IV. APPLICATIONS

### A. State-Space Model of Learning

We begin by demonstrating how the ADMM framework can be applied to a problem with a highly non-linear multimodal measurement model. In the state-space model of learning [28], the system model is a traditional state-space Gauss-Markov process, where the state represents an unobservable cognitive state that represents a subject’s ability to perform a task over time. The corresponding measurement model provides a statistical relationship between the underlying state and the observed task performance for a given trial.

We define  $\mathbf{X} \in \mathbb{R}^{1 \times N}$  to be the cognitive state (with  $K = 1$ ), where  $N$  represents the number of trials conducted. The system model is given by:

$$X_n = \kappa X_{n-1} + \gamma + V_n \quad (19)$$

where  $\kappa \in [0, 1]$  is a forgetting factor,  $\gamma \in \mathbb{R}_+$  is a positive bias that represents a tendency for the cognitive state to in-

crease with time, and  $V_n \sim \mathcal{N}(0, \sigma_V^2)$  is noise in the system model.

Using the state-space model of learning pertaining with multiple behavioral and neurophysiological measures, we assume that each of the  $N$  trials has an associated binary success/failure outcome, a reaction time, and neural spiking behavior. As such, each observation is given by a triplet  $\mathbf{Y}_n = (B_n, R_n, \mathbf{S}_n) \in \{0, 1\} \times \mathbb{R} \times \{0, 1\}^J$ , where  $B_n$  is a binary random variable indicating whether or not the trial was completed successfully,  $R_n$  is the log of the subject's reaction time to complete the task, and  $\mathbf{S}_n$  is a length  $J$  point process that indicates whether or not there was neural spiking activity in each discrete  $\Delta t$  time window.

Each of the three observation modalities is associated with an appropriate statistical model. First, the binary success/failure outcomes obey a Bernoulli probability model:

$$\mathbb{P}(B_n = b_n \mid X_n = x_n) = p_n^{b_n} (1 - p_n)^{1-b_n} \quad (20)$$

where  $p_n$  is given by a logistic function that maps the cognitive state between 0 and 1:

$$p_n = \frac{\exp(\nu + \eta x_n)}{1 + \exp(\nu + \eta x_n)} \quad (21)$$

where  $\nu, \eta \in \mathbb{R}$  are model parameters.

Next, the reaction time obeys a log-normal probability model, with:

$$R_n \sim \mathcal{N}(\psi + \omega X_n, \sigma_R^2) \quad (22)$$

where  $\psi \in \mathbb{R}$  is the estimated initial log reaction time,  $\omega \in \mathbb{R}_-$  is negative to ensure that the reaction time tends to decrease with an increasing cognitive state and  $\sigma_R^2$  represents the level of stochasticity in the relationship between the cognitive state and reaction time.

Lastly, the neural spiking activity is modeled as a point process (as in equation 2.6 of [29]), with the negative log-probability of a given set of spikes given by:

$$\begin{aligned} -\log \mathbb{P}(\mathbf{S}_n = \mathbf{s}_n \mid X_n = x_n) \\ = \sum_{j=1}^J -\log(\Lambda_{n,j}) s_{n,j} + \Lambda_{n,j} \Delta t \end{aligned} \quad (23)$$

where  $s_{n,j} \in \{0, 1\}$  is the  $j$ th bit of  $\mathbf{s}_n$  and  $\log \Lambda$  is the conditional intensity function, given by a generalized linear model [30]:

$$\log \Lambda_{n,j} = \xi + ax_n + \sum_{m=1}^M c_m s_{n,j-m} \quad (24)$$

where  $\xi \in \mathbb{R}$  gives a base intensity level,  $a \in \mathbb{R}$  determines the effect of the cognitive state on the spiking intensity, and  $\mathbf{c} = (c_1, \dots, c_M) \in \mathbb{R}^M$  accounts for the refractory period in neural spiking, i.e., the fact that it is unlikely to see spiking activity in neighboring bins. The point process model given by (23) represents a discrete approximation of the negative log-likelihood for an inhomogeneous Poisson process where the rate in trial  $n$  and time  $j$  is  $\Lambda_{n,j}$ .

Next we adapt the state-space model of learning to the ADMM framework. We begin by considering the negative log-

likelihood of the observations given the underlying cognitive state. We note that not only are the observations temporally conditionally independent given a sequence of cognitive states, but each of the three observations within a trial is conditionally independent given the cognitive state corresponding with that trial:

$$\begin{aligned} L(\mathbf{y} \mid \mathbf{x}) &= \sum_{n=1}^N L_n(\mathbf{y}_n \mid x_n) \\ &= \sum_{n=1}^N L_{B_n}(b_n \mid x_n) + L_{R_n}(r_n \mid x_n) + L_{\mathbf{S}_n}(\mathbf{s}_n \mid x_n) \end{aligned} \quad (25)$$

where the negative log-likelihoods  $L_{B_n} := -\log p_{B_n \mid X_n}$ ,  $L_{R_n} := -\log f_{R_n \mid X_n}$ , and  $L_{\mathbf{S}_n} := -\log p_{\mathbf{S}_n \mid X_n}$  are defined to be the negative log of the appropriate pdf/pmf corresponding with the respective observations. It is important to note that  $L$  is indeed convex. Considering this is not immediately obvious, it is shown in Appendix D.

Next we consider the system model. By defining  $W_n = X_n - \kappa X_{n-1} = \gamma + V_n$  with  $W_0 = X_0$ , we get that  $W_n \sim \mathcal{N}(\gamma, \sigma_V^2)$ , i.e., each  $W_n$  is distributed iid Gaussian. Thus, our negative log-prior is given by:

$$\begin{aligned} \phi(\mathbf{w}) &= -\log \prod_{n=1}^N \mathcal{N}(w_n; \gamma, \sigma_V^2) \\ &\propto \sum_{n=1}^N \frac{(w_n - \gamma)^2}{2\sigma_V^2} \end{aligned} \quad (26)$$

where  $\mathcal{N}(x; \mu, \sigma^2)$  gives the value of a normal distribution with mean  $\mu$  and variance  $\sigma^2$  evaluated at  $x$ . Additionally, under this definition of  $\mathbf{W}$  we get that the transition matrix  $\mathbf{D}$  is in fact just a scalar, namely  $\kappa \in \mathbb{R}$ .

Plugging  $L$ ,  $\phi$ , and  $\mathcal{A}$  into equations (12), (14), and (15), we obtain the update equations for solving the state-space model of learning problem. Beginning with the measurement model update, as a result of its separability across trials, each update decomposes into  $N$  univariate convex minimization problems. As such, these  $N$  problems can be solved in parallel using a convex solver such as CVX [31]. For the system model update, we note that because (14) is separable over  $n = 1, \dots, N$ , the update is reduced to  $N$  quadratic minimizations that can be solved in closed form. Given that the density for  $\mathbf{W}$  is assumed to be fully known, we set the tuning parameter  $\beta = 1$ . The details of these updates can be found in Appendix C.

We demonstrate the state-space model of learning solution on simulated data with  $N = 25$ , using parameters from Section V-A of [21]. The proposed method is compared with the fixed-interval smoother (FIS) detailed in [21] and a sequential Monte Carlo (SMC) method. In particular, we develop a particle smoother using the forward-filtering backward-sampling technique with systematic resampling at each step [32]. For the ADMM method, we set  $\rho = 30$  and limit the procedure to 25 iterations, i.e.,  $\hat{\mathbf{x}} := \mathbf{x}^{(25)}$ . For the SMC method, we use 100 particles. In Table III we look at the average root-mean-square




Fig. 2. Sample realization (blue) for Gaussian state-space model (left) and sparse-variation state-space model (right), along with the estimates using ADMM (red), FIS (green), and SMC (purple). While the Gaussian states are well estimated by all three methods, the ADMM approach utilizing the  $\ell_1$  prior yields the only estimate that captures the piecewise constant nature of the sparse-variation states.

TABLE III  
PERFORMANCE METRICS FOR THE PROPOSED METHOD (ADMM), FIXED-INTERVAL SMOOTHER (FIS), AND SEQUENTIAL MONTE CARLO (SMC) AVERAGED OVER 50 TRIALS WITH THE GAUSSIAN STATE-SPACE MODEL GIVEN BY (19) AND THE STATE-SPACE MODEL WITH SPARSE VARIATIONS GIVEN BY (27)

	Gaussian State		Sparse Variations	
	RMSE	Run Time (s)	RMSE	Run Time (s)
ADMM	<b>0.165</b>	<b>1.8</b>	<b>0.141</b>	7.0
FIS	0.168	2.6	0.181	<b>5.2</b>
SMC	0.188	53.5	0.186	105.7

error (RMSE) and average runtime for each method over 50 trials, where for a given realization  $\mathbf{x}$  and a given estimate  $\hat{\mathbf{x}}$ ,  $\text{RMSE}(\hat{\mathbf{x}}) = \|\hat{\mathbf{x}} - \mathbf{x}\|_2 / \sqrt{N}$ . We note that the proposed method is both most efficient and most accurate in the RMSE sense. While the SMC method would presumably benefit from a larger number of particles, we see that even with limited samples, it is very computationally intensive. While the difference in RMSE is negligible across all 3 methods, it is worth noting that each method obtains a fundamentally different estimate. To be specific, the proposed method gives the MAP estimate in the limit of large iterations, while the other methods yield conditional expectations of the states given the entire observation sequence. In the case of the FIS, the estimate is the conditional expectation under a Gaussian approximation of the posterior. The SMC method, on the other hand, yields the true conditional expectation in the limit of large particle count.

It should be noted that in the case of a Gaussian state space, the problem formulations given by (8) and (9) are nearly equivalent. In particular, it is possible to omit the consensus variable and modify the constraint such that  $\mathbf{W} = \mathbf{X}$ . In such a scenario, the measurement model update would remain the same and the system model update would be solvable with a Kalman smoother. Thus, we further demonstrate the utility of our method by considering a second state-space model with sparse variations where such an approach is not possible. We simulate a state-space model with sparse variations by defining  $X_n = X_{n-1} + V_n$  with  $V_n$  obeying a commonly used sparsity inducing mixture model [33]:

$$V_n = \begin{cases} 0 & \text{w.p. } p \\ \sigma U_n & \text{w.p. } 1 - p \end{cases} \quad (27)$$

where  $p \in [0, 1]$  is a probability,  $\sigma \in \mathbb{R}_+$  is a positive constant, and we define  $U_n \sim \chi_2^2$  as i.i.d. Chi-Squared random variables with two degrees of freedom. This model represents a scenario supported by neurophysiological findings [34], [35] wherein infrequent, discontinuous changes in neural activity arise.

We again conduct 50 trials, setting  $N = 50$ ,  $p = 0.9$ , and  $\sigma = 0.1$ , and estimate the state using ADMM, FIS, and SMC approaches. For the ADMM approach, we note that the true system model is no longer log-concave, so we instead use a sparsity inducing  $\ell_1$  regularizer, i.e., we define  $\phi(\mathbf{w}) = \beta \|\mathbf{w}\|_1$ . As such, we set  $\beta = 15$ , noting that is no longer determined by the model and must be treated as a tuning parameter. The resulting system model update is given by:

$$\mathbf{w}^{(i+1)} = \underset{\mathbf{w}}{\text{argmin}} \frac{\rho}{2} \|\tilde{\mathbf{w}}^{(i)} - \mathbf{w}\|_2^2 + \beta \|\mathbf{w}\|_1.$$

This problem is known as the LASSO problem and may be efficiently solved by applying a soft threshold operation to  $\mathbf{w}^{(i)}$  at each iteration [36].

Given the model mismatch, we observe that the proposed method takes longer to converge on a desirable estimate, and thus increase the maximum number of iterations to 75. For the FIS, given that there is no systematic approach to obtain an estimate with sparse variations, we again utilize a Gaussian approximation, with the noise at each step being modeled by a Gaussian distribution with zero-mean and variance  $\text{Var}(U_n) = 4\sigma^2$ . The SMC method is given the benefit of using the true underlying system model when generating samples on the forward pass. However, when performing the backward pass on sample  $x_n^i$  with respect to a fixed  $\hat{x}_{n+1}$ , we get that when  $x_n^i > \hat{x}_{n+1}$ , the likelihood  $f_{X_{n+1}, \mathbf{Y}|X_n}(\hat{x}_{n+1}, \mathbf{y} | x_n^i) = 0$ , causing the smoother to continually lower  $\hat{x}_n$  for  $n = N, N-1, \dots, 1$  until the smoother fails (i.e.,  $x_k^i > \hat{x}_{k+1}$  for all  $i$  for some  $k \in \{1, \dots, N\}$ ). As such, we only utilize the forward pass particle filter. Referring to Table III for results, we note that the proposed method again outperforms the other methods in the RMSE sense. From a computational perspective, the 3X increase in iterations causes the ADMM approach to take slightly longer than the FIS, though both remain significantly more efficient than the SMC method.

### B. Spectrotemporal Pursuit

Next we demonstrate application of the ADMM framework to the method of spectrotemporal pursuit, originally presented



in [11]. Spectrotemporal pursuit formulates the problem of estimating time varying frequency coefficients as a compressive sensing problem. We define  $\mathbf{Y} \in \mathbb{R}^{P \times N}$  to be a matrix version of an observed time series of length  $PN$ , where each column of  $\mathbf{Y}$  gives a length  $P$  window of the time series. Next, we define  $\mathbf{X} \in \mathbb{R}^{K \times N}$  to be a matrix of frequency coefficients, with each column  $\mathbf{X}_n \in \mathbb{R}^K$  representing the frequency coefficients corresponding with the time window  $\mathbf{Y}_n \in \mathbb{R}^P$ . By defining  $\mathbf{X}$  to be real valued, it is implied that the frequency coefficients are in rectangular form, and thus a frequency resolution of  $K/2$  is achieved. Using this representation, we define the quadratic measurement model:

$$L(\mathbf{y} | \mathbf{x}) = \sum_{n=1}^N \|\mathbf{y}_n - \mathbf{F}_n \mathbf{x}_n\|_2^2 \quad (28)$$

where  $\mathbf{F}_n \in \mathbb{R}^{P \times K}$  is an inverse Fourier matrix, i.e.,  $(\mathbf{F}_n)_{p,k} := \cos(2\pi((n-1)P+p)\frac{k-1}{K})$  and  $(\mathbf{F}_n)_{p,k+\frac{K}{2}} := \sin(2\pi((n-1)P+p)\frac{k-1+K/2}{K})$  for  $p = 1, \dots, P$  and  $k = 1, \dots, K/2$ . In this sense we can view the spectrotemporal estimation problem as a traditional linear measurement with Gaussian noise problem. As such, it is well defined when  $P \geq K$ , which is consistent with the well known fact that the number of frequency coefficients associated with a time series can not exceed the number of samples.

The method of spectrotemporal pursuit removes this constraint by introducing a sparsity inducing prior on the frequency coefficients, paralleling the approaches in compressive sensing used to estimate the coefficients underlying a system with an underdetermined set of observations. In particular, spectrotemporal pursuit imposes a group-sparsity prior on the first differences of the frequency coefficients. Letting  $\mathbf{W}_n = \mathbf{X}_n - \mathbf{X}_{n-1}$  (i.e.,  $\mathbf{D}$  is the identity matrix), we define the system model:

$$\phi(\mathbf{w}) = \sum_{k=1}^K \left( \sum_{n=1}^N w_{k,n}^2 \right)^{\frac{1}{2}}. \quad (29)$$

We can view this function as the  $\ell_1$ -norm of a vector whose entries are the  $\ell_2$ -norms of the rows of the argument. As such,  $\phi(\mathbf{w})$  is small when only a small number of the rows of  $\mathbf{w}$  are non-zero. Furthermore, the rows that are non-zero should have a small  $\ell_2$ -norm. Application of this function to the differences of the frequency coefficients over time ensures that throughout a given time series, most frequency coefficients do not vary, and those that do vary are varying smoothly. This time-frequency characterization is known to occur in certain biological time-series. Thus, spectrotemporal pursuit utilizes this knowledge to obtain significantly denoised spectrotemporal estimates while avoiding the time/frequency resolution trade-off without necessitating a sliding window approach. This is again reminiscent of compressive sensing, which makes strong claims regarding the recoverability of a set of coefficients with underdetermined measurements so long as the coefficients are sufficiently sparse.

The spectrotemporal pursuit solution initially proposed in [11] is an iteratively reweighted least squares (IRLS) algorithm. While the IRLS algorithm is also exact and offers convergence guarantees, it requires inversion of  $N \times N$  and  $K \times K$  matrices

$N$  times per iteration of the algorithm. Furthermore, design of the state-covariance matrix obfuscates the problem and requires careful thought when modifying the system model.

The proposed ADMM framework yields a straightforward solution to the spectrotemporal pursuit problem. First, plugging  $L$  into equation (12) yields:

$$\begin{aligned} \mathbf{x}_n^{(i+1)} &= \underset{\mathbf{x}_n}{\operatorname{argmin}} \|\mathbf{y}_n - \mathbf{F}_n \mathbf{x}_n\|_2^2 + \frac{\rho}{2} \|\mathbf{x}_n - \tilde{\mathbf{x}}_n^{(i)}\|_2^2 \\ &= \underset{\mathbf{x}_n}{\operatorname{argmin}} \|\mathbf{x}_n + \mathbf{C}_n \mathbf{b}_n^{(i)}\|_{\mathbf{C}_n}^2 \\ &= -\mathbf{C}_n \mathbf{b}_n^{(i)} \end{aligned} \quad (30)$$

where  $\mathbf{C}_n := (\mathbf{F}_n^T \mathbf{F}_n + \frac{\rho}{2} \mathbf{I})^{-1}$  and  $\mathbf{b}_n^{(i)} := -\frac{1}{2}(\mathbf{F}_n^T \mathbf{y}_n + \rho \tilde{\mathbf{x}}_n^{(i)})$ . We note that when  $P < K$ ,  $\mathbf{F}_n^T \mathbf{F}_n$  is rank deficient and it is our choice of  $\rho$  that ensures the update is well formed. Also, it is important to note that each  $\mathbf{C}_n$  for  $n = 1, \dots, N$  can be computed once at initialization, as they do not change throughout iterations.

Next, placing the group-sparsity prior in equation (14) shows that the system model update is given by a standard group-lasso problem:

$$\mathbf{w}^{(i+1)} = \underset{\mathbf{w}}{\operatorname{argmin}} \|\tilde{\mathbf{w}}^{(i)} - \mathbf{w}\|_2^2 + \frac{2\beta}{\rho} \sum_{k=1}^K \left( \sum_{n=1}^N w_{k,n}^2 \right)^{\frac{1}{2}}. \quad (31)$$

Furthermore, this special case with an orthonormal regressor matrix (i.e., the identity) yields a closed form solution, namely a row-wise shrinkage operator applied to  $\tilde{\mathbf{w}}^{(i)}$  [24]. The shrinkage amount is proportional to the tuning parameter  $\beta$ , with larger  $\beta$  yielding a smaller number of non-zero rows in  $\mathbf{w}$ .

We demonstrate the ADMM solution for spectrotemporal pursuit on a simulated example recreated from the original paper [11]. Let  $\tilde{\mathbf{y}} \in \mathbb{R}^M$  be the vectorized version of  $\mathbf{y}$  with  $M = NP$  and  $\mathbf{y}_n = [\tilde{y}_{(n-1)P+1}, \tilde{y}_{(n-1)P+2}, \dots, \tilde{y}_{nP}]^T$  for  $n = 1, \dots, N$ . Then, we consider the signal:

$$\begin{aligned} \tilde{y}_m &= 10 \cos^8(2\pi f_0 m) \sin(2\pi f_1 m) \\ &\quad + 10 \exp\left(4\frac{m-M}{M}\right) \cos(2\pi f_2 m) + v_m \end{aligned} \quad (32)$$

where  $f_0 = 0.04$  Hz,  $f_1 = 10$  Hz,  $f_2 = 11$  Hz, and  $v_m \sim \mathcal{N}(0, 1)$  iid for  $m = 1, \dots, M$ . Letting the sampling frequency be  $f_s = 125$  Hz and  $M = 7500$  gives a simulated time-series 600 seconds in duration. We note that  $\mathbf{y}$  contains a sparse number of active frequency components, and the frequency components that are active are modulated over time in a smooth fashion. Additionally, the active frequency components  $f_1$  and  $f_2$  are chosen to be in neighboring frequencies, creating an increased difficulty when trying to distinguish their respective contributions.

The top row of Fig. 3 shows time-frequency estimates of the simulated time-series using traditional methods and spectrotemporal pursuit. First, we observe that the standard spectrogram (Fig. 3 A) suffers from significant spectral leakage and is unable to clearly distinguish between the 10 Hz and 11 Hz frequency components. For the spectrotemporal pursuit estimate (Fig. 3 B)

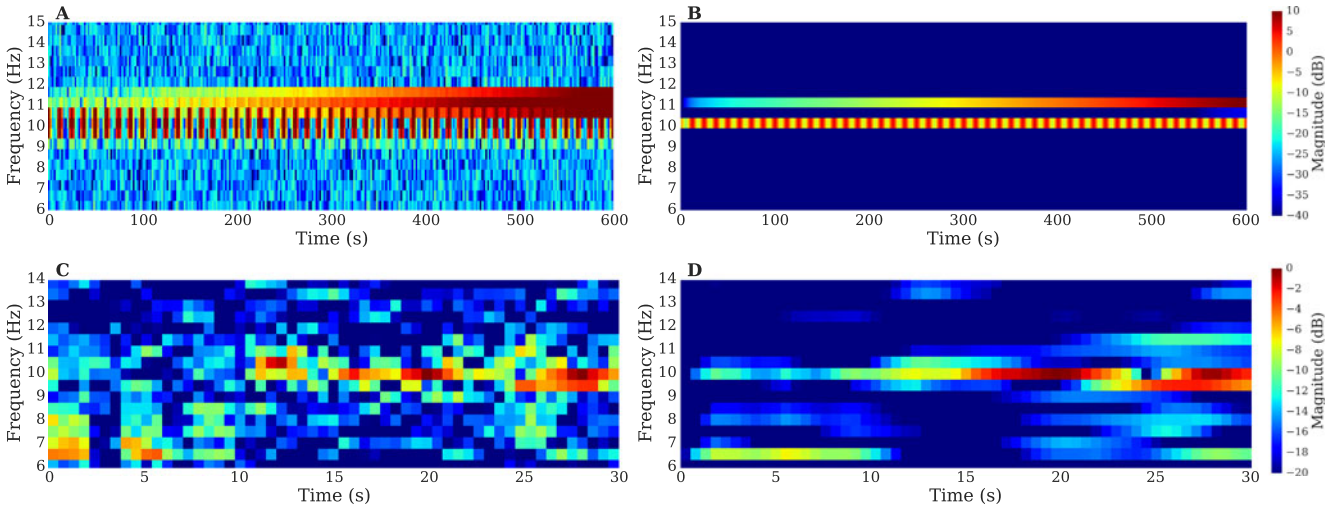


Fig. 3. Spectrotemporal decompositions for simulated time series given by (32) (A/B) and single channel EEG recording (C/D). **A**: Traditional spectrogram with  $NFFT = 2f_s$ , no overlap, and Hanning window. **B**: Spectrotemporal pursuit estimate with  $K = 2f_s$ ,  $P = K/8$ . **C**: Traditional spectrogram with  $NFFT = 1024$ , 75% overlap and Hanning window. **D**: Low-Rank Spectrotemporal Decomposition with  $K = 1024$  and  $P = K/4$ .

we select  $P < K$ , meaning that the number of samples in each time window is less than the number of frequency bins. As such, we are effectively increasing the temporal resolution while still maintaining the spectral resolution without the use of overlapping windows. Because this would in general be an underdetermined problem, the group-sparsity prior is needed to ensure the problem has a unique solution. In addition to increased temporal resolution, we witness that spectrotemporal pursuit enables the contributions from  $f_1$  and  $f_2$  to be clearly distinguishable. Further benefits of this approach to spectrotemporal decompositions are given in detail in [11]. Here, we are proposing an algorithm that offers improvements in efficiency, modularity, and interpretability. In particular, we witness a roughly  $10\times$  speedup per iteration on the same size data when using the ADMM framework rather than IRLS.

To further illustrate the modularity of the proposed framework, we next demonstrate that we can utilize an entirely different system model with a minor adjustment to a single update. Specifically, we consider a low-rank spectrotemporal decomposition (LRSD) which substitutes the nuclear norm for the group sparsity prior [37]. As such, the LRSD estimate is obtained by substituting the system model update given by (31) with:

$$\mathbf{w}^{(i+1)} = \underset{\mathbf{w}}{\operatorname{argmin}} \|\tilde{\mathbf{w}}^{(i)} - \mathbf{w}\|_F^2 + \beta \|\mathbf{w}\|_* \quad (33)$$

where  $\|\cdot\|_*$  is the nuclear norm, given by the sum of the singular values of the argument. Conveniently, this update is known as the matrix lasso and yields a straightforward solution via singular value soft thresholding [38]. By making a simple adjustment to the means by which  $\mathbf{w}^{(i)}$  is updated, we are able to obtain an entirely different spectrotemporal decomposition.

This point is illustrated by the bottom row of Fig. 3 where we demonstrate the LRSD on human single-channel EEG data using adhesive flexible sensors [39]. The data in question contains a 30-second recording in which the subject’s eyes are closed at the 10 second mark, at which point we would expect to see

increased energy in the alpha band (10–12 Hz). The change point nature of the recording suggests that the group sparsity prior on the dynamics, which enforces smoothness across time, is ill-suited for this recording, and the traditional spectrogram (Fig. 3 C) suffers significantly from noise. By not explicitly enforcing smoothness in time, the low-rank enforcing nuclear norm prior (Fig. 3 D) accommodates the change point and is able to significantly suppress activity outside of the alpha band. Similarly to the spectrotemporal pursuit example, we are able to set  $P < K$  and achieve equivalent temporal resolution to the spectrogram without utilizing overlapping windows or sacrificing spectral resolution.

*Remark 2:* Comparisons with other methods are intentionally omitted in this section given that there is no systematic application to these non-Markov problem formulation. While the original problem proposed in equation (7) does not lend itself to an obvious solution for the discussed non-Markov models, the consensus ADMM formulation given by (11) may be solved in a straightforward manner. In particular, we note that the EKF and UKF have no clear extensions for non-Markov scenarios and the use of sampling based methods for such models would require drawing samples of group-sparse or low-rank matrices.

## V. DISCUSSION

We have presented a unified framework for solving a broad class of dynamic modeling problems. The proposed method can be applied to systems with non-linear measurements and/or non-Markov dynamics. As demonstrated on two applications, our framework can be applied in a straightforward manner to acquire efficient solutions to problems that may otherwise require complex or approximate solutions. Furthermore, we have shown that this algorithm will converge on the true MAP estimate of the latent signal in the limit of large iterations. With this provably accurate algorithm comes a mathematical justification for an intuitive approach to dynamic time-series es-

timization, namely iteratively computing estimates based on the measurement model and system model and then averaging them in the appropriate sense.

There are a number of extensions to this framework still to be explored. The most glaring shortcomings are the inability to conduct the estimation procedure causally and the necessity to know model parameters a priori. Regarding the former, we note the use of homotopy schemes for causal estimation that gradually incorporate new observations into the solution [40], [41]. Additionally, there has been recent research investigating algorithms for performing ADMM in an online fashion [42], [43] that could potentially be leveraged by our framework. To address the latter, expectation-maximization (EM) techniques can be built into the ADMM iterations in order to estimate model parameters jointly with the desired latent time-series. In that regard, the E-step, which requires sampling from the posterior distribution, is typically the bottleneck. To address that, Langevin based methods and stochastic gradient descent methods can be used to efficiently sample from the posterior distribution [44]. Identifying sufficient conditions on mixing times for generating approximately i.i.d. posterior samples for the M-step could be the subject of future in-depth work. We note that while there exist sample based methods for estimating model parameters [19, Sec. IV], these methods can be computationally prohibitive as witnessed in Table III.

Lastly, we note that there is considerable interest in state-space estimation where the observations or system are subject to noise from heavy-tailed distributions such as the Student's  $t$  or Cauchy distributions [45], [46], which are not log-concave. Recent literature has shown that ADMM can be shown to converge under even milder conditions than those assumed by Theorem 1 [47], [48]. Given that both the Student's  $t$  and Cauchy distributions are log-quasi-concave, continuous, and possess a single local maximum, we could reasonably expect convergence of our framework to the MAP estimate in such a scenario. This topic provides interesting opportunities for future experimental and theoretical work.

#### APPENDIX A DERIVATION OF SCALED FORM

We will demonstrate the derivation of the scaled form of the measurement model update only, noting that the derivation for the other updates follows almost identical steps. Consider the original measurement model update:

$$\begin{aligned} \mathbf{x}^{(i+1)} = \operatorname{argmin}_{\mathbf{x}} & \left( \sum_{n=1}^N L_n(\mathbf{y}_n | \mathbf{x}_n) \right) + \langle \lambda^{(i)}, \mathbf{x} - \mathbf{z}^{(i)} \rangle \\ & + \frac{\rho}{2} \left\| \mathbf{x} - \mathbf{z}^{(i)} \right\|_F^2. \end{aligned} \quad (34)$$

For ease of notation, the superscript  $(i)$  is omitted for the remainder of this appendix. Using the definition of the inner product and Frobenius norm, we can break up the second and third terms

across into sums and simplify as follows:

$$\begin{aligned} \mathbf{x}^{(i+1)} = \operatorname{argmin}_{\mathbf{x}} & \sum_{n=1}^N L_n(\mathbf{y}_n | \mathbf{x}_n) + \lambda_n^T (\mathbf{x}_n - \mathbf{z}_n) \\ & + \frac{\rho}{2} (\mathbf{x}_n - \mathbf{z}_n)^T (\mathbf{x}_n - \mathbf{z}_n) \\ = \operatorname{argmin}_{\mathbf{x}} & \sum_{n=1}^N L_n(\mathbf{y}_n | \mathbf{x}_n) + \frac{\rho}{2} \mathbf{x}_n^T \mathbf{x}_n \\ & + (\lambda_n - \rho \mathbf{z}_n)^T \mathbf{x}_n \\ = \operatorname{argmin}_{\mathbf{x}} & \sum_{n=1}^N \frac{2}{\rho} L_n(\mathbf{y}_n | \mathbf{x}_n) + \mathbf{x}_n^T \mathbf{x}_n \\ & - 2 \left( \mathbf{z}_n - \frac{\lambda_n}{\rho} \right)^T \mathbf{x}_n. \end{aligned} \quad (35)$$

Defining  $\tilde{\mathbf{x}}_n = \mathbf{z}_n - \frac{\lambda_n}{\rho}$  as in Section III-A, we note that  $\tilde{\mathbf{x}}_n$  does not depend on  $\mathbf{x}$ , enabling us to complete the square and simplify as follows:

$$\begin{aligned} \mathbf{x}_n = \operatorname{argmin}_{\mathbf{x}} & \sum_{n=1}^N \frac{2}{\rho} L_n(\mathbf{y}_n | \mathbf{x}_n) + \mathbf{x}_n^T \mathbf{x}_n - 2\tilde{\mathbf{x}}_n^T \mathbf{x}_n \\ = \operatorname{argmin}_{\mathbf{x}} & \sum_{n=1}^N \frac{2}{\rho} L_n(\mathbf{y}_n | \mathbf{x}_n) + \mathbf{x}_n^T \mathbf{x}_n - 2\tilde{\mathbf{x}}_n^T \mathbf{x}_n + \tilde{\mathbf{x}}_n^T \tilde{\mathbf{x}}_n \\ = \operatorname{argmin}_{\mathbf{x}} & \sum_{n=1}^N \frac{2}{\rho} L_n(\mathbf{y}_n | \mathbf{x}_n) + (\tilde{\mathbf{x}}_n - \mathbf{x}_n)^T (\tilde{\mathbf{x}}_n - \mathbf{x}_n) \\ = \operatorname{argmin}_{\mathbf{x}} & \left( \sum_{n=1}^N L_n(\mathbf{y}_n | \mathbf{x}_n) \right) + \frac{\rho}{2} \|\mathbf{x} - \tilde{\mathbf{x}}\|_F^2, \end{aligned} \quad (36)$$

as was to be shown.

#### APPENDIX B PROOF OF THEOREM 1

Consider the problem in its original form:

$$\begin{aligned} (\hat{\mathbf{x}}, \hat{\mathbf{w}}) = \operatorname{argmin}_{\mathbf{x}, \mathbf{w}} & L(\mathbf{y} | \mathbf{x}) + \beta \phi(\mathbf{w}) \\ \text{s.t.} & \mathbf{w} = \mathcal{A}(\mathbf{x}). \end{aligned} \quad (37)$$

The goal is to show that there is an equivalent two-block ADMM problem whose updates match those given by (11). To do so we define the variable  $\mathbf{Q} := [\mathbf{X}^T, \mathbf{W}^T]^T \in \mathbb{R}^{2K \times N}$  ( $\mathbf{X}, \mathbf{W} \in \mathbb{R}^{K \times N}$ ) and the function  $g(\mathbf{Q}) := L(\mathbf{y} | \mathbf{X}) + \beta \phi(\mathbf{W})$ . Next, we define  $\mathbf{Z} := [\mathbf{Z}_X^T, \mathbf{Z}_W^T]^T \in \mathbb{R}^{2K \times N}$  and the function:

$$h(\mathbf{Z}) = \begin{cases} 0 & \mathcal{A}(\mathbf{Z}_X) = \mathbf{Z}_W \\ \infty & \mathcal{A}(\mathbf{Z}_X) \neq \mathbf{Z}_W \end{cases}. \quad (38)$$

Using these newly defined terms, we can write (37) equivalently as:

$$\begin{aligned} (\hat{\mathbf{q}}, \hat{\mathbf{z}}) &= \underset{\mathbf{q}, \mathbf{z}}{\operatorname{argmin}} \quad g(\mathbf{q}) + h(\mathbf{z}) \\ \text{s.t.} \quad & \mathbf{q} - \mathbf{z} = 0. \end{aligned} \quad (39)$$

Note that if  $\mathbf{q} = [\mathbf{x}^T, \mathbf{w}^T]^T$  is such that  $\mathbf{w} \neq \mathcal{A}(\mathbf{x})$  (the constraints in (37) are not satisfied) and  $\mathbf{z}$  is such that  $\mathbf{q} - \mathbf{z} = 0$  (the constraints (39) are satisfied), then  $h(\mathbf{z}) = \infty$  and  $(\mathbf{q}, \mathbf{z})$  are not the minimizers of (39). To solve this problem with ADMM, we first find augmented Lagrangian:

$$\mathcal{L}_\rho(\mathbf{q}, \mathbf{z}, \boldsymbol{\gamma}) = g(\mathbf{q}) + h(\mathbf{z}) + \langle \boldsymbol{\gamma}, \mathbf{q} - \mathbf{z} \rangle + \frac{\rho}{2} \|\mathbf{q} - \mathbf{z}\|_F^2 \quad (40)$$

with Lagrange multiplier  $\boldsymbol{\gamma} = [\lambda^T, \boldsymbol{\alpha}^T]^T \in \mathbb{R}^{2K \times N}$  ( $\lambda, \boldsymbol{\alpha} \in \mathbb{R}^{K \times N}$ ). As a result, we get the following update equations:

$$\begin{aligned} \mathbf{q}^{(i+1)} &= \underset{\mathbf{q}}{\operatorname{argmin}} \mathcal{L}_\rho(\mathbf{q}, \mathbf{z}^{(i)}, \boldsymbol{\gamma}^{(i)}) \\ \mathbf{z}^{(i+1)} &= \underset{\mathbf{z}}{\operatorname{argmin}} \mathcal{L}_\rho(\mathbf{q}^{(i+1)}, \mathbf{z}, \boldsymbol{\gamma}^{(i)}) \\ \boldsymbol{\gamma}^{(i+1)} &= \boldsymbol{\gamma}^{(i)} + \rho(\mathbf{q}^{(i+1)} - \mathbf{z}^{(i+1)}). \end{aligned} \quad (41)$$

Next we show that the update equations given by (41) are equivalent to those given by (11).

First, consider the  $\mathbf{q}$  update:

$$\begin{aligned} \mathbf{q}^{(i+1)} &= \underset{\mathbf{q}}{\operatorname{argmin}} \mathcal{L}_\rho(\mathbf{q}, \mathbf{z}^{(i)}, \boldsymbol{\gamma}^{(i)}) \\ &= \underset{\mathbf{q}}{\operatorname{argmin}} g(\mathbf{q}) + \langle \boldsymbol{\gamma}^{(i)}, \mathbf{q} - \mathbf{z}^{(i)} \rangle + \frac{\rho}{2} \|\mathbf{q} - \mathbf{z}^{(i)}\|_F^2 \\ &= \underset{[\mathbf{x}^T, \mathbf{w}^T]^T}{\operatorname{argmin}} L(\mathbf{y} | \mathbf{x}) + \beta\phi(\mathbf{w}) \\ &\quad + \left[ \begin{array}{c} \lambda^{(i)} \\ \boldsymbol{\alpha}^{(i)} \end{array} \right]^T \left( \left[ \begin{array}{c} \mathbf{x} \\ \mathbf{w} \end{array} \right] - \left[ \begin{array}{c} \mathbf{z}_x^{(i)} \\ \mathbf{z}_w^{(i)} \end{array} \right] \right) \\ &\quad + \left\| \left[ \begin{array}{c} \mathbf{x} \\ \mathbf{w} \end{array} \right] - \left[ \begin{array}{c} \mathbf{z}_x^{(i)} \\ \mathbf{z}_w^{(i)} \end{array} \right] \right\|_F^2 \\ &= \left[ \begin{array}{c} \mathbf{x}^{(i+1)} \\ \mathbf{w}^{(i+1)} \end{array} \right] \end{aligned} \quad (42)$$

where  $\mathbf{x}^{(i+1)}$  and  $\mathbf{w}^{(i+1)}$  are given by:

$$\mathbf{x}^{(i+1)} = \underset{\mathbf{x}}{\operatorname{argmin}} L(\mathbf{y} | \mathbf{x}) + \langle \lambda^{(i)}, \mathbf{x} - \mathbf{z}_x^{(i)} \rangle + \|\mathbf{x} - \mathbf{z}_x^{(i)}\|_F^2 \quad (43)$$

$$\mathbf{w}^{(i+1)} = \underset{\mathbf{w}}{\operatorname{argmin}} \beta\phi(\mathbf{w}) + \langle \boldsymbol{\alpha}^{(i)}, \mathbf{w} - \mathbf{z}_w^{(i)} \rangle + \|\mathbf{w} - \mathbf{z}_w^{(i)}\|_F^2 \quad (44)$$

and can be found independently of each other.

Next, consider the  $\mathbf{z}$  update:

$$\begin{aligned} \mathbf{z}^{(i+1)} &= \underset{\mathbf{z}}{\operatorname{argmin}} \mathcal{L}_\rho(\mathbf{q}^{(i+1)}, \mathbf{z}, \boldsymbol{\gamma}^{(i)}) \\ &= \underset{\mathbf{z}}{\operatorname{argmin}} h(\mathbf{z}) + \langle \boldsymbol{\gamma}^{(i)}, \mathbf{q}^{(i+1)} - \mathbf{z} \rangle \\ &\quad + \frac{\rho}{2} \|\mathbf{q}^{(i+1)} - \mathbf{z}\|_F^2 \\ &= \underset{[\mathbf{z}_x^T, \mathbf{z}_w^T]^T}{\operatorname{argmin}} h(\mathbf{z}) + \langle \lambda^{(i)}, \mathbf{x}^{(i+1)} - \mathbf{z}_x \rangle \\ &\quad + \frac{\rho}{2} \|\mathbf{x}^{(i+1)} - \mathbf{z}_x\|_F^2 \\ &\quad + \langle \boldsymbol{\alpha}^{(i)}, \mathbf{w}^{(i+1)} - \mathbf{z}_w \rangle \\ &\quad + \frac{\rho}{2} \|\mathbf{w}^{(i+1)} - \mathbf{z}_w\|_F^2 \\ &= \underset{\mathbf{z}_x}{\operatorname{argmin}} \langle \lambda^{(i)}, \mathbf{x}^{(i+1)} - \mathbf{z}_x \rangle \\ &\quad + \frac{\rho}{2} \|\mathbf{x}^{(i+1)} - \mathbf{z}_x\|_F^2 \\ &\quad + \langle \boldsymbol{\alpha}^{(i)}, \mathbf{w}^{(i+1)} - \mathcal{A}(\mathbf{z}_x) \rangle \\ &\quad + \frac{\rho}{2} \|\mathbf{w}^{(i+1)} - \mathcal{A}(\mathbf{z}_x)\|_F^2 \\ &= \left[ \begin{array}{c} \mathbf{z}_x^{(i+1)} \\ \mathcal{A}(\mathbf{z}_x^{(i+1)}) \end{array} \right] \end{aligned}$$

where  $\mathbf{z}_x^{(i+1)}$  is given as the solution to (15), i.e., the consensus update for our target problem, and the second to last equality follows from the fact that  $h(\mathbf{z})$  is infinite if  $\mathbf{z}_w \neq \mathcal{A}(\mathbf{z}_x)$ , so we can treat the problem as a single variable optimization problem.

Next we can substitute these results into the equations for the  $\mathbf{q}$  update to obtain:

$$\begin{aligned} \mathbf{w}^{(i+1)} &= \underset{\mathbf{w}}{\operatorname{argmin}} \beta\phi(\mathbf{w}) + \langle \boldsymbol{\alpha}^{(i)}, \mathbf{w} - \mathcal{A}(\mathbf{z}_x^{(i)}) \rangle \\ &\quad + \|\mathbf{w} - \mathcal{A}(\mathbf{z}_x^{(i)})\|_F^2 \end{aligned} \quad (45)$$

which is the (unscaled) update equation (14) for  $\mathbf{w}$  in the original formulation, where  $\mathbf{z}_x^{(i)}$  in this formulation corresponds with  $\mathbf{z}^{(i)}$  in the original formulation. The  $\mathbf{x}$  portion of the  $\mathbf{q}$  remains unchanged from (43), which is equivalent to the unscaled update equation (12) for  $\mathbf{x}$  in the original formulation.

Next, we can decompose the matrix multiplication in the same way as above to show that:

$$\boldsymbol{\gamma}^{(i+1)} = \left[ \begin{array}{c} \lambda^{(i+1)} \\ \boldsymbol{\alpha}^{(i+1)} \end{array} \right] \quad (46)$$

where  $\lambda^{(i+1)}$  and  $\boldsymbol{\alpha}^{(i+1)}$  are given by the original updates in (11).

Thus, we have shown that directly solving (39) using ADMM yields the proposed updates detailed in the body of the paper. As such, we will show that the ADMM solution to (39) is convergent. By assumption,  $L$  and  $\phi$  are closed, proper, and convex, and hence, so is their sum  $g$ . To show that  $h$  is convex, we note

that this is an indicator function on the set  $H := \{(\mathbf{z}_X, \mathbf{z}_W) : \mathcal{A}(\mathbf{z}_X) = \mathbf{z}_W\} \subset \mathbb{R}^2$ , thus  $h$  is convex if and only if  $H$  is convex [49, Ch. 2]. Suppose  $\mathbf{z}^1 = [\mathbf{z}_X^1, \mathbf{z}_W^1]^T$  and  $\mathbf{z}^2 = [\mathbf{z}_X^2, \mathbf{z}_W^2]^T$  are such that  $\mathcal{A}(\mathbf{z}_X^1) = \mathbf{z}_W^1$  and  $\mathcal{A}(\mathbf{z}_X^2) = \mathbf{z}_W^2$ , i.e.,  $\mathbf{z}^1, \mathbf{z}^2 \in H$ . Then, if we take a convex combination  $\mathbf{z}^\alpha := \alpha \mathbf{z}^1 + (1 - \alpha) \mathbf{z}^2$  for  $\alpha \in [0, 1]$ , we get:

$$\begin{aligned} \mathbf{z}_W^\alpha &= \alpha \mathbf{z}_W^1 + (1 - \alpha) \mathbf{z}_W^2 \\ &= \alpha \mathcal{A}(\mathbf{z}_X^1) + (1 - \alpha) \mathcal{A}(\mathbf{z}_X^2) \\ &= \mathcal{A}(\alpha \mathbf{z}_X^1 + (1 - \alpha) \mathbf{z}_X^2) \\ &= \mathcal{A}(\mathbf{z}_X^\alpha). \end{aligned} \quad (47)$$

Thus, we see that  $\mathbf{z}^1, \mathbf{z}^2 \in H \Rightarrow \mathbf{z}^\alpha \in H$ , i.e.,  $H$ , and therefore  $h$ , are convex. It then follows from Section 3.2.1 of [15] that the ADMM solution for (39) is convergent, as was to be shown. ■

### APPENDIX C STATE-SPACE MODEL OF LEARNING UPDATES

We begin by deriving expressions for the negative log-likelihoods for each of the observations:

$$\begin{aligned} L_{B_n}(b_n | x_n) &= -\log p_{B_n | X_n}(b_n | x_n) \\ &= -\log p_n^{b_n} (1 - p_n)^{1-b_n} \\ &= -b_n \log \frac{e^{\nu + \eta x_n}}{1 + e^{\nu + \eta x_n}} - (1 - b_n) \log \frac{1}{1 + e^{\nu + \eta x_n}} \\ &\propto \log(1 + e^{\nu + \eta x_n}) - b_n \eta x_n \end{aligned}$$

$$\begin{aligned} L_{R_n}(r_n | x_n) &= -\log f_{R_n | X_n}(r_n | x_n) \\ &= -\log \frac{1}{\sqrt{2\pi\sigma_R^2}} \exp\left(-\frac{(r_n - \psi - \omega x_n)^2}{2\sigma_R^2}\right) \\ &\propto \frac{(r_n - \psi - \omega x_n)^2}{2\sigma_R^2} \end{aligned}$$

$$\begin{aligned} L_{S_n}(s_n | x_n) &= -\log p_{S_n | X_n}(s_n | x_n) \\ &= -\log \exp\left(\sum_{j=1}^J [\log(\Lambda_{n,j}) s_{n,j} - \Lambda_{n,j} \Delta t]\right) \\ &= -\sum_{j=1}^J \left(\xi + ax_n + \sum_{m=1}^M c_m s_{n,j-m}\right) n_{n,j} \\ &\quad + \sum_{j=1}^J \exp\left(\xi + ax_n + \sum_{m=1}^M c_m s_{n,j-m}\right) \Delta t \\ &\propto -ax_n \sum_{j=1}^J n_{n,j} \\ &\quad + \sum_{j=1}^J \exp\left(\xi + ax_n + \sum_{m=1}^M c_m s_{n,j-m}\right) \Delta t \end{aligned}$$

$$\begin{aligned} &= \Delta t \exp(\xi + ax_n) \sum_{j=1}^J \exp\left(\sum_{m=1}^M c_m s_{n,j-m}\right) \\ &\quad - ax_n \sum_{j=1}^J n_{n,j} \end{aligned}$$

These expressions can be plugged into equation (25) to obtain the measurement model update equation, which can in turn be solved using Newton's method.

Next, the system model update can be solved in closed form:

$$\begin{aligned} \mathbf{w}^{(i+1)} &= \underset{\mathbf{w}}{\operatorname{argmin}} \phi(\mathbf{w}) + \frac{\rho}{2} \|\mathbf{w} - \tilde{\mathbf{w}}^{(i)}\|^2 \\ &= \underset{\mathbf{w}}{\operatorname{argmin}} \sum_{n=1}^N \left( \frac{(w_n - \gamma)^2}{2\sigma_V^2} + \frac{\rho}{2} (w_n - \tilde{w}_n^{(i)})^2 \right) \end{aligned}$$

where  $\tilde{w}_n^{(i)} := z_n^{(i)} - \kappa z_{n-1}^{(i)} - \alpha_n^{(i)} / \rho$ . Thus, we can solve for each  $w_n$  separately:

$$\begin{aligned} w_n^{(i+1)} &= \underset{w_n}{\operatorname{argmin}} \frac{(w_n - \gamma)^2}{2\sigma_V^2} + \frac{\rho}{2} (w_n - \tilde{w}_n^{(i)})^2 \\ &= \underset{w_n}{\operatorname{argmin}} \left( \frac{1}{2\sigma_V^2} + \frac{\rho}{2} \right) w_n^2 - \left( \frac{\gamma}{\sigma_V^2} + \rho \tilde{w}_n^{(i)} \right) w_n \\ &= \underset{w_n}{\operatorname{argmin}} \left( w_n - \frac{\frac{\gamma}{\sigma_V^2} + \rho \tilde{w}_n^{(i)}}{\frac{1}{\sigma_V^2} + \rho} \right)^2 \\ &= \frac{\frac{\gamma}{\sigma_V^2} + \rho \tilde{w}_n^{(i)}}{\frac{1}{\sigma_V^2} + \rho}. \end{aligned}$$

Finally, given its relatively low dimensionality, we can efficiently solve the consensus update in closed form by posing it as a least squares problem. First, we note that  $\mathcal{A}(\mathbf{z}) = \mathbf{G}\mathbf{z}$  when we define:

$$\mathbf{G} = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 & 0 \\ -\kappa & 1 & 0 & \dots & 0 & 0 \\ 0 & -\kappa & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & 0 \\ 0 & 0 & 0 & \dots & -\kappa & 1 \end{bmatrix} \quad (48)$$

with  $\mathbf{G} \in \mathbb{R}^{N \times N}$ . Thus, we have:

$$\mathbf{z}^{(i+1)} = \underset{\mathbf{z}}{\operatorname{argmin}} \|\mathbf{z} - \tilde{\mathbf{z}}_X^{(i)}\|_F^2 + \|\mathbf{G}\mathbf{z} - \tilde{\mathbf{z}}_W^{(i)}\|_F^2. \quad (49)$$

Taking the gradient of the RHS and setting to zero yields:

$$\mathbf{z}^{(i+1)} = (\mathbf{I} + \mathbf{G}^T \mathbf{G})^{-1} (\tilde{\mathbf{z}}_X^{(i)} + \mathbf{G}^T \tilde{\mathbf{z}}_W^{(i)}). \quad (50)$$

Given that  $\mathbf{G}$  is known a-priori, we can find  $(\mathbf{I} + \mathbf{G}^T \mathbf{G})^{-1}$  once and each consensus update becomes a matrix multiplication problem.

## APPENDIX D

## CONVEXITY STATE-SPACE MODEL OF LEARNING NEGATIVE LOG-LIKELIHOOD

Given that  $L$  is the sum of the negative log-likelihoods for each of the observation modalities as in (25), it is sufficient to show that they are each convex in  $x_n$ , which is made easier by use of the simplifications derived in Appendix C. Noting that addition of a constant does not affect convexity, we can assess the final simplification provided in each case. As such, we see that  $L_{B_n}(b_n | x_n)$  is the sum of a term that is linear in  $x_n$  and a special case of the log sum exponential (LSE) function with an added auxiliary variable constrained to equal zero (giving  $e^0 = 1$ ). Given the convexity of LSE, its sum with a linear term is also convex, and thus  $L_{B_n}(b_n | x_n)$  is convex. Next,  $L_{R_n}(r_n | x_n)$  is quadratic in  $x_n$  and thus convex. Finally,  $L_{S_n}(s_n | x_n)$  is the sum of a term that is linear in  $x_n$  and a term that is exponential in  $x_n$ , both of which are convex. As a result,  $L_{B_n}(b_n | x_n)$ ,  $L_{R_n}(r_n | x_n)$ , and  $L_{S_n}(s_n | x_n)$  are all convex in  $x_n$  for any  $(b_n, r_n, s_n) \in \{0, 1\} \times \mathbb{R} \times \{0, 1\}^J$ , and thus so is their sum  $L$ .

## ACKNOWLEDGMENT

The authors would like to thank M. Wagner for his contributions to the development of the LRSD algorithm and the reviewers for their insightful comments and suggestions for improvements.

## REFERENCES


- [1] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, Feb. 2002.
- [2] W. D. Penny, K. J. Friston, J. T. Ashburner, S. J. Kiebel, and T. E. Nichols, *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. New York, NY, USA: Academic, 2011.
- [3] R. Dugad and N. Ahuja, "Video denoising by combining Kalman and Wiener estimates," in *Proc. IEEE Int. Conf. Image Process.*, 1999, vol. 4, pp. 152–156.
- [4] R. E. Kalman, "A new approach to linear filtering and prediction problems," *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, 1960.
- [5] E. A. Wan and R. Van Der Merwe, "The unscented Kalman filter for nonlinear estimation," in *Proc. IEEE Adapt. Syst. Signal Process., Commun., Control Symp.*, 2000, pp. 153–158.
- [6] P. Del Moral, "Non-linear filtering: Interacting particle resolution," *Markov Process. Related Fields*, vol. 2, no. 4, pp. 555–581, 1996.
- [7] N. Vaswani, "Kalman filtered compressed sensing," in *Proc. 15th IEEE Int. Conf. Image Process.*, 2008, pp. 893–896.
- [8] D. Angelosante, G. B. Giannakis, and E. Grossi, "Compressed sensing of time-varying signals," in *Proc. 16th Int. Conf. Digit. Signal Process.*, 2009, pp. 1–8.
- [9] J. Ziniel, L. C. Potter, and P. Schniter, "Tracking and smoothing of time-varying sparse signals via approximate belief propagation," in *Proc. Conf. Rec. 44th Asilomar Conf. Signals, Syst. Comput.*, 2010, pp. 808–812.
- [10] A. S. Charles, A. Balavoine, and C. J. Rozell, "Dynamic filtering of time-varying sparse signals via L1 minimization," *IEEE Trans. Signal Process.*, vol. 64, no. 21, pp. 5644–5656, Nov. 2015.
- [11] D. Ba, B. Babadi, P. L. Purdon, and E. N. Brown, "Robust spectrotemporal decomposition by iteratively reweighted least squares," *Proc. Nat. Acad. Sci.*, vol. 111, no. 50, pp. E5336–E5345, 2014.
- [12] A. Charles, M. S. Asif, J. Romberg, and C. Rozell, "Sparsity penalties in dynamical system estimation," in *Proc. 45th Annu. Conf. Inf. Sci. Syst.*, 2011, pp. 1–6.
- [13] D. Angelosante, S. I. Roumeliotis, and G. B. Giannakis, "Lasso-Kalman smoother for tracking sparse signals," in *Proc. Conf. Record 43rd Asilomar Conf. Signals, Syst. Comput.*, 2009, pp. 181–185.
- [14] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical Learning With Sparsity*. Boca Raton, FL, USA: CRC Press, 2015.
- [15] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
- [16] M. Sahraee-Ardakan and A. K. Fletcher, "Estimation and learning of dynamic nonlinear networks (DyNNets)," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2017, pp. 2856–2860.
- [17] F. Pérez and B. E. Granger, "IPython: A system for interactive scientific computing," *Comput. Sci. Eng.*, vol. 9, no. 3, pp. 21–29, May 2007. [Online]. Available: <http://ipython.org>
- [18] S. J. Godsill, A. Doucet, and M. West, "Monte Carlo smoothing for nonlinear time series," *J. Amer. Statist. Assoc.*, vol. 99, no. 465, pp. 156–168, 2004.
- [19] O. Cappé, S. J. Godsill, and E. Moulines, "An overview of existing methods and recent advances in sequential Monte Carlo," *Proc. IEEE*, vol. 95, no. 5, pp. 899–924, May 2007.
- [20] C. Andrieu, A. Doucet, and R. Holenstein, "Particle Markov chain Monte Carlo methods," *J. Roy. Statist. Soc., B (Statist. Methodol.)*, vol. 72, no. 3, pp. 269–342, 2010.
- [21] T. P. Coleman, M. Yanike, W. A. Suzuki, and E. N. Brown, "A mixed-filter algorithm for dynamically tracking learning from multiple behavioral and neurophysiological measures," *The Dynamic Brain: An Exploration of Neuronal Variability and its Functional Significance*. London, U.K.: Oxford Univ. Press, 2011.
- [22] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [23] N. Parikh and S. P. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, pp. 127–239, 2014.
- [24] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight, "Sparsity and smoothness via the fused Lasso," *J. Roy. Statist. Soc., B (Statist. Methodol.)*, vol. 67, no. 1, pp. 91–108, 2005.
- [25] G. H. Golub and C. F. Van Loan, *Matrix Computations*, vol. 3. Baltimore, MD, USA: JHU Press, 2012.
- [26] C. Chen, B. He, Y. Ye, and X. Yuan, "The direct extension of ADMM for multi-block convex minimization problems is not necessarily convergent," *Math. Program.*, vol. 155, no. 1/2, pp. 57–79, 2016.
- [27] J. Eckstein and W. Yao, "Understanding the convergence of the alternating direction method of multipliers: Theoretical and computational perspectives," *Pac. J. Optim.*, vol. 11, pp. 619–644, 2015.
- [28] A. C. Smith *et al.*, "Dynamic analysis of learning in behavioral experiments," *J. Neurosci.*, vol. 24, no. 2, pp. 447–461, 2004.
- [29] T. P. Coleman and S. S. Sarma, "A computationally efficient method for nonparametric modeling of neural spiking activity with point processes," *Neural Comput.*, vol. 22, no. 8, pp. 2002–2030, 2010.
- [30] W. Truccolo, U. T. Eden, M. R. Fellows, J. P. Donoghue, and E. N. Brown, "A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects," *J. Neurophysiol.*, vol. 93, no. 2, pp. 1074–1089, 2005.
- [31] CVX Research, "CVX: Matlab software for disciplined convex programming, version 2.0," Aug. 2012. [Online]. Available: <http://cvxr.com/cvx>
- [32] A. Doucet and A. M. Johansen, "A tutorial on particle filtering and smoothing: Fifteen years later," *Handbook Nonlinear Filtering*, vol. 12, pp. 656–704, 2009.
- [33] Y. Wu and S. Verdú, "Optimal phase transitions in compressed sensing," *IEEE Trans. Inf. Theory*, vol. 58, no. 10, pp. 6241–6263, Oct. 2012.
- [34] L. D. Lewis *et al.*, "Rapid fragmentation of neuronal networks at the onset of propofol-induced unconsciousness," *Proc. Nat. Acad. Sci.*, vol. 109, no. 49, pp. E3377–E3386, 2012.
- [35] D. Ba, B. Babadi, P. L. Purdon, and E. N. Brown, "Neural spike train denoising by point process re-weighted iterative smoothing," in *Proc. 48th Asilomar Conf. Signals, Syst. Comput.*, 2014, pp. 763–768.
- [36] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Roy. Statist. Soc. B (Methodological)*, vol. 58, pp. 267–288, 1996.
- [37] G. Schamberg, D. Ba, M. Wagner, and T. Coleman, "Efficient low-rank spectrotemporal decomposition using ADMM," in *Proc. IEEE Statist. Signal Process. Workshop*, 2016, pp. 1–5.
- [38] S. Ma, D. Goldfarb, and L. Chen, "Fixed point and Bregman iterative methods for matrix rank minimization," *Math. Program.*, vol. 128, no. 1, pp. 321–353, 2011.
- [39] D. Y. Kang, Y.-S. Kim, G. Ornelas, M. Sinha, K. Naidu, and T. P. Coleman, "Scalable microfabrication procedures for adhesive-integrated flexible and stretchable electronic sensors," *Sensors*, vol. 15, no. 9, pp. 23 459–23 476, 2015.

- [40] M. S. Asif, A. Charles, J. Romberg, and C. Rozell, "Estimation and dynamic updating of time-varying signals with sparse variations," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2011, pp. 3908–3911.
- [41] M. S. Asif and J. Romberg, "Dynamic updating for sparse time varying signals," in *Proc. 43rd Annu. Conf. Inf. Sci. Syst.*, 2009, pp. 3–8.
- [42] H. Wang and A. Banerjee, "Online alternating direction method (longer version)," arXiv preprint arXiv:1306.3721, 2013.
- [43] S. Liu, J. Chen, P.-Y. Chen, and A. O. Hero, "Zeroth-order online alternating direction method of multipliers: Convergence analysis and applications," arXiv:1710.07804, 2017.
- [44] S. Mandt, M. D. Hoffman, and D. M. Blei, "Stochastic gradient descent as approximate Bayesian inference," arXiv:1704.04289, 2017.
- [45] G. Agamennoni, J. I. Nieto, and E. M. Nebot, "Approximate inference in state-space models with heavy-tailed noise," *IEEE Trans. Signal Process.*, vol. 60, no. 10, pp. 5024–5037, Oct. 2012.
- [46] Y. Huang, Y. Zhang, Z. Wu, N. Li, and J. Chambers, "A novel robust students t-based Kalman filter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 53, no. 3, pp. 1545–1554, Jun. 2017.
- [47] Y. Wang, W. Yin, and J. Zeng, "Global convergence of ADMM in non-convex nonsmooth optimization," arXiv:1511.06324, 2015.
- [48] S. Magnússon, P. C. Weeraddana, M. G. Rabbat, and C. Fischione, "On the convergence of alternating direction Lagrangian methods for nonconvex structured optimization problems," *IEEE Trans. Control Netw. Syst.*, vol. 3, no. 3, pp. 296–309, Sep. 2016.
- [49] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ, USA: Princeton Univ. Press, 2015.



**Gabriel Schamberg** (S'16) received the B.S. degree in computer engineering in 2012 and the M.S. degree in electrical engineering in 2016, both from the University of California, San Diego (UCSD), San Diego, CA, USA, where he is currently working toward the Ph.D. degree in electrical engineering. From 2012 to 2014, he was a Software Developer with NKI Engineering, San Diego, CA, USA. In 2014, he began his graduate studies and was the recipient of the Jacobs Fellowship from the Department of Electrical and Computer Engineering, UCSD. Since 2014, he

has been with the Neural Interaction Lab at UCSD.



**Demba Ba** (M'18) received the B.Sc. degree in electrical engineering from the University of Maryland, College Park, MD, USA, in 2004, and the M.Sci. and Ph.D. degrees in electrical engineering and computer science with a minor in mathematics from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2006 and 2011, respectively. In 2006 and 2009, he was a Summer Research Intern with the Communication and Collaboration Systems Group, Microsoft Research, Redmond, WA, USA. From 2011 to 2014, he was a Postdoctoral Associate with the MIT/Harvard Neuroscience Statistics Research Laboratory, where he developed theory and efficient algorithms to assess synchrony among large assemblies of neurons. He is currently an Assistant Professor of electrical engineering and bioengineering with Harvard University, where he directs the CRISP group. His research interests lie at the intersection of high-dimensional statistics, optimization, dynamic modeling, and computational neuroscience. In 2016, he was the recipient of a Research Fellowship in Neuroscience from the Alfred P. Sloan Foundation.



**Todd P. Coleman** (SM'11) received the Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2005. He was a Postdoctoral Scholar in neuroscience with MIT and Massachusetts General Hospital during the 2005–2006 academic year. He was an Assistant Professor of the electrical and computer engineering as well as neuroscience with the University of Illinois from 2006 to 2011. He is currently a Professor of bioengineering, a Principal Investigator of the Neural Interaction Laboratory, and the

Codirector of the Center for Perinatal Health with the University of California San Diego, La Jolla, CA, USA. His research interests lie at the intersection of applied mathematics, bioelectronics, neuroscience, and medicine. He was named Gilbreth Lecturer for the National Academy of Engineering in 2015.