



# Spatio-Temporal Dynamics of Intrinsic Networks in Functional Magnetic Imaging Data Using Recurrent Neural Networks

R. Devon Hjelm<sup>1,2\*</sup>, Eswar Damaraju<sup>3,4†</sup>, Kyunghyun Cho<sup>5</sup>, Helmut Laufs<sup>6</sup>, Sergey M. Plis<sup>3</sup> and Vince D. Calhoun<sup>3,4</sup>

<sup>1</sup> Montréal Institute for Learning Algorithms, Montreal, QC, Canada, <sup>2</sup> Microsoft Research, Montreal, QC, Canada, <sup>3</sup> The Mind Research Network, Albuquerque, NM, United States, <sup>4</sup> The University of New Mexico, Albuquerque, NM, United States, <sup>5</sup> New York University, New York, NY, United States, <sup>6</sup> Schleswig University Hospital, Kiel, Germany

## OPEN ACCESS

### Edited by:

Anqi Qiu,  
National University of Singapore,  
Singapore

### Reviewed by:

Chao-Gan Yan,  
Chinese Academy of Sciences, China  
Jong-Hwan Lee,  
Korea University, South Korea

### \*Correspondence:

R. Devon Hjelm  
eroneus@gmail.com

†These authors have contributed  
equally to this work

### Specialty section:

This article was submitted to  
Brain Imaging Methods,  
a section of the journal  
Frontiers in Neuroscience

Received: 23 January 2018

Accepted: 09 August 2018

Published: 20 September 2018

### Citation:

Hjelm RD, Damaraju E, Cho K,  
Laufs H, Plis SM and Calhoun VD  
(2018) Spatio-Temporal Dynamics of  
Intrinsic Networks in Functional  
Magnetic Imaging Data Using  
Recurrent Neural Networks.  
Front. Neurosci. 12:600.  
doi: 10.3389/fnins.2018.00600

We introduce a novel recurrent neural network (RNN) approach to account for temporal dynamics and dependencies in brain networks observed via functional magnetic resonance imaging (fMRI). Our approach directly parameterizes temporal dynamics through recurrent connections, which can be used to formulate blind source separation with a conditional (rather than marginal) independence assumption, which we call RNN-ICA. This formulation enables us to visualize the temporal dynamics of both first order (activity) and second order (directed connectivity) information in brain networks that are widely studied in a static sense, but not well-characterized dynamically. RNN-ICA predicts dynamics directly from the recurrent states of the RNN in both task and resting state fMRI. Our results show both task-related and group-differentiating directed connectivity.

**Keywords:** deep learning, ICA, RNN, neuroimaging methods, fMRI, resting-state fMRI, AOD

## 1. INTRODUCTION

Functional magnetic resonance imaging (fMRI) of blood oxygenation-level dependent (BOLD) signal provides a powerful tool for studying temporally coherent patterns in the brain (Damoiseaux et al., 2006; Calhoun et al., 2008; Smith et al., 2009). Intrinsic networks (INs, Biswal et al., 1995) and functional connectivity are important outcomes of fMRI studies which illuminate our understanding of healthy and diseased brain function (Calhoun et al., 2001b; Allen et al., 2012). While deep or non-linear approaches for INs from fMRI and MRI exist (Hjelm et al., 2014; Plis et al., 2014; Castro et al., 2016), of the tools available, the most widely used are *generative models* with shallow and linear structure. Such models typically use a shared parameterization of structure to learn a common model across subjects which refactor the data into a constrained space that both provides straightforward analysis and allows for efficient and effective learning algorithms.

The most popular of such methods, independent component analysis (ICA, Bell and Sejnowski, 1995), begins with the hypothesis that the data is a mixture of maximally independent sources. ICA is trainable through one of many relatively simple optimization routines that maximize non-Gaussianity or minimize mutual information (Hyvärinen and Oja, 2000). However, ICA, as with other popular linear methods for separating INs, is order-agnostic in time: each multivariate signal at each time step is treated as independent and identically distributed (i.i.d.). While model

degeneracy in time is convenient for learning; as an assumption about the data the explicit lack of temporal dependence necessarily marginalizes out dynamics, which then must be extrapolated in *post-hoc* analysis.

In addition, ICA, as it is commonly used in fMRI studies, uses the same parameterization across subjects, which allows for either temporal or spatial variability, but not both (Calhoun et al., 2001a). The consequence of this is that ICA is not optimized to represent variation of shape in INs while also representing variation in time courses. This may encourage ICA to exaggerate time course statistics, as any significant variability in shape or size will primarily be accounted for by the time courses.

Despite these drawbacks, the benefits of using ICA for separating independent sources in fMRI data is strongly evident in numerous studies, to the extent that has become the dominant approach for separating INs and analyzing connectivity (Damoiseaux et al., 2006; Calhoun et al., 2008; Kim et al., 2008; Smith et al., 2009; Zuo et al., 2010; Allen et al., 2012; Calhoun and Adali, 2012). In order to overcome shortcomings in temporal dynamics and subject/temporal variability, but without abandoning the fundamental strengths of ICA, we extend ICA to model sequences using recurrent neural networks (RNNs). The resulting model, which we call RNN-ICA, naturally represents temporal dynamics through a sequential ICA objective and is easily trainable using back-propagation and gradient descent.

## 2. BACKGROUND

Here we will formalize the problem of source separation with temporal dependencies and formulate the solution in terms of maximum likelihood estimation (MLE) and a recurrent model that parameterizes a conditionally independent distribution (i.e., RNNs).

Let us assume that the data is composed of  $N$  ordered sequences of length  $T$ ,

$$\mathbf{X}_n = (\mathbf{x}_{1,n}, \mathbf{x}_{2,n}, \dots, \mathbf{x}_{T,n}), \quad (1)$$

where each element in the sequence,  $\mathbf{x}_{t,n}$ , is a  $D$  dimensional vector, and the index  $n$  enumerates the whole sequence. The goal is to find/infer a set of source signals,

$$\mathbf{S}_n = (\mathbf{s}_{1,n}, \mathbf{s}_{2,n}, \dots, \mathbf{s}_{T,n}), \quad (2)$$

such that a subsequence  $\mathbf{s}_{t_1:t_2} = (\mathbf{s}_{t_1,n}, \mathbf{s}_{t_1+1,n}, \dots, \mathbf{s}_{t_2,n})$  generates a subsequence of data,  $\mathbf{x}_{t'_1:t'_2} = (\mathbf{x}_{t'_1,n}, \mathbf{x}_{t'_1+1,n}, \dots, \mathbf{x}_{t'_2,n})$ , for  $t_1 \leq t'_1 < t'_2$  and  $t_1 < t_2 \leq t'_2$ . In particular, we are interested in finding a generating function,

$$\mathbf{X}_n = G(\mathbf{S}_n, \epsilon), \quad (3)$$

where  $\epsilon$  is an additional noise variable.

This problem can generally be understood as *inference* of unobserved or latent configurations from time-series observations. It is convenient to assume that the sources,  $\mathbf{S}_n$ , are stochastic random variables with well-understood and

interpretable noise, such as Gaussian or logistic variables with independence constraints. Representable as a directed graphical model in time, the choice of a-priori model structure, such as the relationship between latent variables and observations, can have consequences on model capacity and inference complexity.

Directed graphical models often require complex approximate inference which introduces variance into learning. Rather than solving the general problem in Equation (3), we will assume that the generating function,  $G(\cdot)$ , is *noiseless*, and the source sequences,  $\mathbf{S}_n$  have the same dimensionality as the data,  $\mathbf{X}_n$ , with each source signal being composed of a set of conditionally independent components with density parameterized by a RNN. We will show that the learning objective closely resembles that of noiseless independent component analysis (ICA). Assuming generation is noiseless and preserves dimensionality will reduce variance which would otherwise hinder learning with high-dimensional, low-sample size data, such as fMRI.

### 2.1. Independent Component Analysis

ICA (Bell and Sejnowski, 1995) hypothesizes that the observed data is a linear mixture of independent sources:  $\mathbf{x}_{t,n} = \sum_m s_{t,n,m} \mathbf{m}_m$ , where  $\mathbf{s}_{t,n} = \{s_{t,n,m}\}$  are sources and  $\mathbf{m}_m$  are the columns of a mixing matrix,  $\mathbf{M}$ . ICA constrains the sources (a.k.a., *components*) to be maximally independent. This framework presupposes any specific definition of component independence, and algorithms widely used for fMRI typically fall under two primary families, kurtosis-based methods and infomax (Hyvärinen and Oja, 2000), although there are other algorithms providing a more flexible density estimation (Fu et al., 2014).

For the infomax algorithm (Bell and Sejnowski, 1995), the model is parameterized by an unmixing matrix  $\mathbf{W} = \mathbf{M}^{-1}$ , such that  $\mathbf{S}_n = f(\mathbf{X}_n) = \mathbf{W} \cdot \mathbf{X}_n$ . In the context of fMRI, the infomax objective seeks to minimize the mutual information of  $\mathbf{s}_{n,t}$  for all subjects at all times. This can be shown to be equivalent to assuming the prior density of the sources are non-Gaussian and that they factorize, or  $p_s(\mathbf{s}_{t,n}) = \prod_{m=1}^M p_{s_m}(s_{t,n,m})$ , where  $\mathbf{s}_{t,n} = \{s_{t,n,m}\}$  is an  $M$ -dimensional vector. When the sources are drawn from a logistic distribution, it can be shown that infomax is equivalent to MLE, with the log-likelihood objective for the empirical density,  $p_x(\mathbf{X}_n)$ , being transformed by  $f(\mathbf{X}) = \mathbf{W} \cdot \mathbf{X}$ :

$$\log p_x(\mathbf{X}_n) = \log p_s(\mathbf{W} \cdot \mathbf{X}_n) + T \log |\det \mathbf{W}|, \quad (4)$$

where  $|\det \mathbf{W}| = |\det \mathbf{J}_f(\mathbf{X})|$  is the absolute value of the determinant of the Jacobian matrix.

With ICA, generating example sequences can be done by applying the inverse of the unmixing matrix to an ordered set of sources. However, one cannot simply sample from the model and generate samples of the observed data: any attempt to do so would simply generate unordered data and not true sequences. The sources in ICA are constrained to be marginally independent in time; ICA does not explicitly model dynamics, and training on shuffled observed sequences will regularly produce the same source structure.

There are numerous graphical models and methods designed to model sequences, including hidden Markov models (HMMs)

and sequential Monte Carlo (SMC, Doucet et al., 2001). HMMs are a popular and simple generative directed graphical models in time with tractable inference and learning and a traditional approach in modeling language. However, HMMs place a high burden on the hidden states to encode enough long-range dynamics to model entire sequences. Recurrent neural networks (RNNs), on the other hand, have the capacity to encode long-range dependencies through deterministic hidden units. When used in conjunction to the ICA objective, the resulting algorithm is a novel and, as we will show, much more powerful, approach for blind source separation based on a conditional independence assumption.

## 2.2. Recurrent Neural Networks

An RNN is a type of neural network with cyclic connections that has seen widespread success in neural machine translation (Cho et al., 2014), sequence-to-sequence learning (Sutskever et al., 2014), sequence generation (Graves, 2013), and numerous other settings. When computing the internal state across a sequence index (such as time or word/character position), RNNs apply the same set of parameters (i.e., connective *weights*) at each step. This gives the model the properties of translational symmetry and directed dependence across time, which are desirable if we expect directed dependence with the same update rules across the sequence. In addition, this makes RNN relatively memory-efficient, as one set of parameters are used across the sequence dimension.

RNNs have many forms, but we will focus on those that act as probabilistic models of sequences, i.e.:

$$p(\mathbf{X}) = p(\mathbf{x}_1:T) = p(\mathbf{x}_1) \prod_{t=2}^T p(\mathbf{x}_t | \mathbf{x}_1:t-1). \quad (5)$$

Better known as a “language model” or *generative RNN*, the exact form of the conditional density typically falls under a family of transformations,

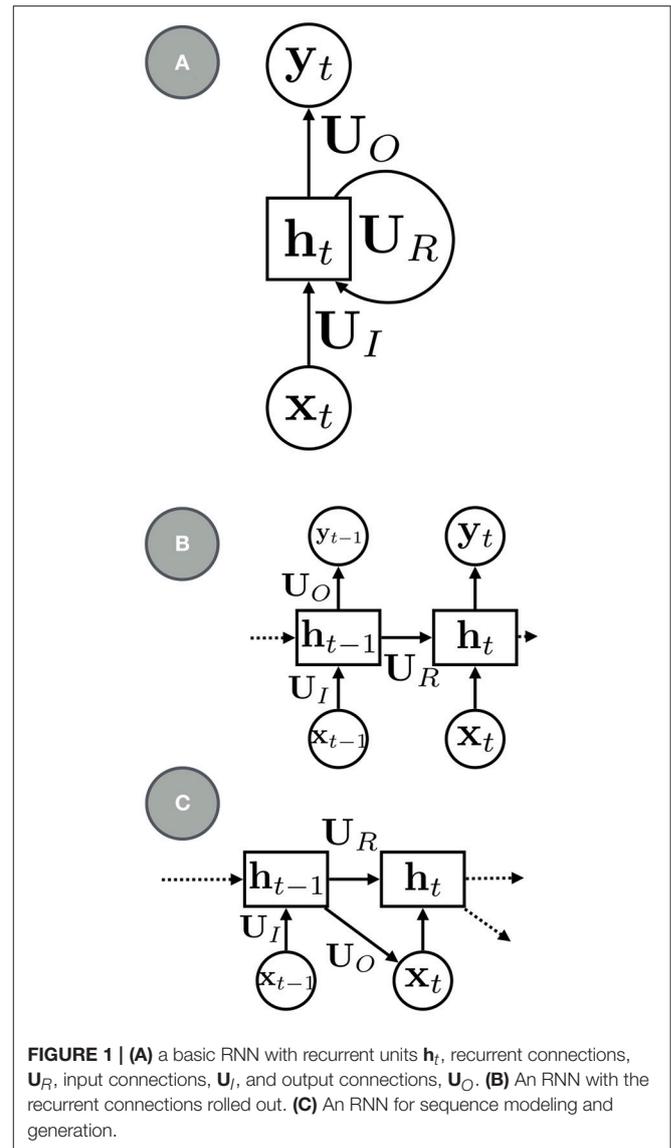
$$p(\mathbf{x}_t | \mathbf{x}_1:t-1) = f(\mathbf{h}_t; \boldsymbol{\psi}); \quad \mathbf{h}_t = g(\mathbf{h}_{t-1}, \mathbf{x}_{t-1}; \boldsymbol{\phi}), \quad (6)$$

where  $\mathbf{h}_t$  are a set of deterministic *recurrent* states (or “recurrent unit”).  $g(\cdot; \boldsymbol{\phi})$  are recurrent connections that take the current observation and hidden state as input and output the next recurrent state. The output connections,  $f(\cdot; \boldsymbol{\psi})$ , take the recurrent states as input at each step and output the parameters for the conditional distribution. Note that the model parameters,  $\boldsymbol{\psi}$  and  $\boldsymbol{\phi}$ , are *recurrent*: the same parameters are used at every time step and are not unique across the sequence index,  $t$ .

The most canonical RNN for sequence modeling has the a simple parameterization (e.g., see **Figure 1**):

$$g(\mathbf{h}_t; \boldsymbol{\psi}) = \tanh(\mathbf{U}_R \mathbf{h}_{t-1} + \mathbf{U}_I \mathbf{x}_{t-1} + \mathbf{b}), \quad (7)$$

where  $\mathbf{U}_R$  is a square matrix of recurrent weights,  $\mathbf{U}_I$  are the input weights, and  $\mathbf{b}$  is a bias term. The mappings between the various variables in the model need not be shallow: an RNN with deep neural networks can model more complex recurrent transitions. Parameterizations that use gating and other types of memory



functions, such as long short-term memory (LSTM, Hochreiter and Schmidhuber, 1997) and gated recurrent units (GRUs, Cho et al., 2014), can be used to better model longer sequences and are also widely used.

Training an RNN for simple sequence modeling is easily done with the back-propagation algorithm, using the negative log-likelihood objective over the output conditional distributions:

$$\mathcal{L} = -\frac{1}{N} \sum_{n=1}^N \left( \sum_{t=2}^T \log p(\mathbf{x}_t | \mathbf{x}_1:t-1) + \log p(\mathbf{x}_1) \right). \quad (8)$$

Typically the loss is computed with mini-batches instead of over the entire dataset for efficiency, randomizing mini-batches at each training epoch. The marginal density,  $p(\mathbf{x}_1)$ , can be learned by fitting to the average marginal across time, either to parameters of a target distribution directly or by training a neural

network to predict the hidden state that generates  $\mathbf{x}_1$  (Bahdanau et al., 2014).

### 3. METHODS

#### 3.1. RNN-ICA

RNNs have already been shown to be capable of predicting signal from BOLD fMRI data (Güçlü and van Gerven, 2017), though usually in the supervised setting. An unsupervised RNN framework for sequence modeling can easily be extended to incorporate the infomax objective. Define, as with ICA, a linear transformation for each observation to source configuration:  $\mathbf{s}_{t,n} = \mathbf{W}\mathbf{x}_{t,n}$ , and define a high-kurtosis and factorized source distribution,  $p_{s_{t,n}}(\mathbf{s}_{t,n})$  (such as a logistic or Laplace distribution) for each time step,  $t$ , and each fMRI sequence,  $n$ . We apply this transformation to an fMRI time series:  $\mathbf{s}_{1:T,n} = f(\mathbf{x}_{1:T,n}) = (\mathbf{W}\mathbf{x}_{1,n}, \mathbf{W}\mathbf{x}_{2,n}, \dots, \mathbf{W}\mathbf{x}_{T,n})$ . The log-likelihood function over the whole sequence,  $\mathbf{X}_n = \mathbf{x}_{1:T,n}$ , can be re-parameterized as:

$$\begin{aligned} \log p(\mathbf{x}_{1:T,n}) &= \log p(\mathbf{x}_{1,n}) + \sum_{t=2}^T \log p(\mathbf{x}_{t,n} | \mathbf{x}_{1:t-1,n}) \\ &= \log p_{s_1}(\mathbf{W}\mathbf{x}_{1,n}) + \sum_{t=2}^T \log p_{s_{t,n}}(\mathbf{W}\mathbf{x}_{t,n} | \mathbf{x}_{1:t-1,n}) \\ &\quad + \log |\det \mathbf{J}_f(\mathbf{x}_{1:T,n})| \\ &= T \log |\det \mathbf{W}| + \log p_{s_{1,n}}(\mathbf{W}\mathbf{x}_{1,n}) + \sum_{t=2}^T \log p_{s_{t,n}} \\ &\quad (\mathbf{W}\mathbf{x}_{t,n} | \mathbf{x}_{1:t-1,n}), \end{aligned} \quad (9)$$

where  $\mathbf{J}_f$  is the Jacobian over the transformation,  $f$ , and the source distribution,  $p_{s_{t,n}}$ , has parameters determined by the recurrent states,  $\mathbf{h}_{t,n}$ . A high-kurtosis distribution is desirable to ensure independence of the sources (or minimizing the mutual information, e.g., the infomax objective Bell and Sejnowski, 1995), so a reasonable choice for the outputs of the RNN at each time step are the mean,  $\boldsymbol{\mu}$ , and scale,  $\boldsymbol{\sigma}$ , for a logistic distribution:

$$\boldsymbol{\mu}_{t,n} = \mathbf{W}_\mu \mathbf{h}_{t,n}; \quad \boldsymbol{\sigma}_{t,n} = \mathbf{W}_\sigma \mathbf{h}_{t,n}. \quad (10)$$

**Figure 2** illustrates the network structure for a few time steps as well as the forward and back-propagated signal, and **Algorithm 1** demonstrates the training procedure for RNN-ICA. For our model, all network parameters and the ICA weight / un-mixing matrix,  $\mathbf{W}$ , are the same for all subjects at all times. Our treatment assumes the ICA weight matrix is square, which is necessary to ensure a tractable determinant Jacobian and inverse. fMRI data is very high dimensional, so to reduce the dimensionality, we must resort to some sort of dimensionality reduction as preprocessing. A widely used for dimensionality reduction in ICA studies of fMRI is principle component analysis (PCA) (Calhoun et al., 2001b; Allen et al., 2012), used to reduce the data to match the selected number of sources,  $\mathbf{s}_{t,n}$ .

Note that RNNs with deeper architectures have been very successful for generative tasks (e.g., WaveNets, Van Den Oord et al., 2016), and RNN-ICA could benefit from a deeper

---

#### Algorithm 1 : RNN-ICA

---

```

W ← initialize the unmixing matrix
θ = (UR, UI, b) ← initialize the RNN recurrent
and input weights and bias (see Equation 7)
ψ = (Wμ, Wσ) ← initialize the RNN output weights (see
Equation 10)
D = {x1:T,n}n=1N ← N sequences of PCA-preprocessed fMRI
sequences, windowed at T time steps.
repeat
  x1:T,n ~ D ▷ Draw random samples from the set of PCA
sequences
  s1:T,n ← (Wx1,n, Wx2,n, ... WxT,n) ▷ Transform the PCA
component sequence with the unmixing matrix
  h1,n ← f(x1,n) ▷ Initialize the first hidden state, as
discussed in Section 2.2
  (μ1,n, σ1,n) ← (Wμh1,n, Wσh1,n) ▷ Compute the
parameters of the initial probability distribution
  for t from 2 to T do
    ht,n ← tanh(URht-1,n + UIxt-1,n + b) ▷ Update each
hidden state and conditional in-sequence
    (μt,n, σt,n) ← (Wμht,n, Wσht,n) ▷ Compute the
parameters of the conditional probability at time t
  end for
  ℒ ← -T log |det W| - log  $p_{s_{1,n}}(\mathbf{s}_{1,n})$  +
 $\sum_{t=2}^T \log p_{s_{t,n}}(\mathbf{s}_{t,n} | \mathbf{x}_{1:t-1,n})$  ▷ Compute the negative log
likelihood
  φ = (W, θ, ψ) ← φ -  $\gamma \nabla_{\phi} \mathcal{L}$  ▷ Perform gradient descent
on the parameters
until convergence

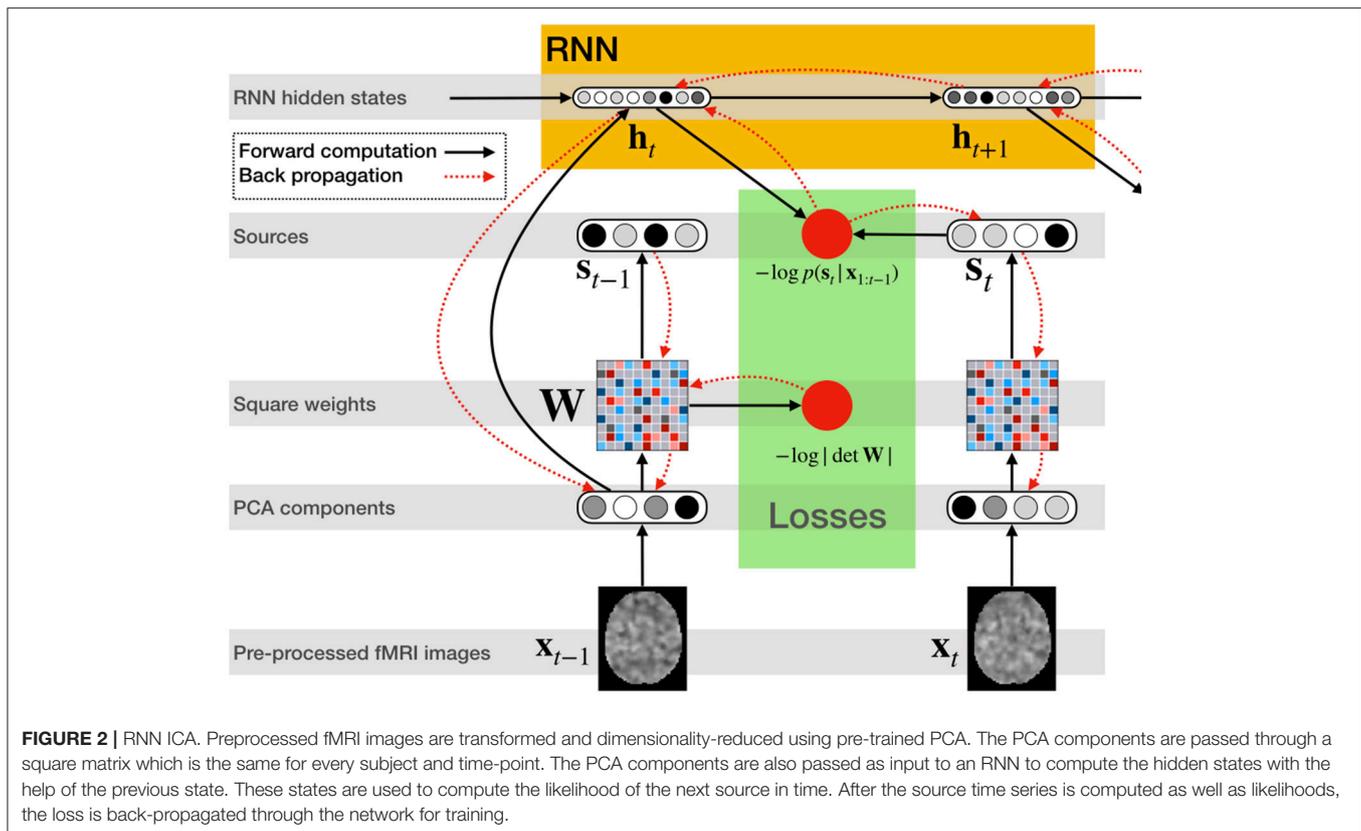
```

---

architecture capable of inferring more complex relationships in the data. However, as fMRI data is often composed of a low number of training samples, we found it necessary to demonstrate the ability of RNN-ICA to learn meaningful sources with a simple RNN architecture. We leave architectural improvements for RNN-ICA for future research.

### 4. EXPERIMENTS AND RESULTS

We first apply RNN-ICA to synthetic data simulated to evaluate the model performance and subsequently on real functional magnetic imaging (fMRI) data. fMRI analyses typically falls under two categories: task-based and resting state analysis. Task experiments typically involve subjects being exposed to a time-series of stimulus, from which task-specific components can be extrapolated. In the case of RNN-ICA, this should reveal task-related directed connectivity and spatial variability, in addition to the usual task-relatedness of activity from ICA. Resting-state data is often used to confirm the presence of distinct and functional states of the brain. We chose a dataset resting state experiment that also had simultaneous Electroencephalography (EEG) from which ground-state subject neurobiological states could be derived. For RNN-ICA, we should be able to find a correspondence between predicted activation as defined in our model and changes in state. As a result, this should



provide a means to prevent false positives or negatives when interpreting resting state network or inter-group differences owing to (systematically) different sleep stages present in their examined cohorts.

#### 4.1. Experiments With Simulated Data

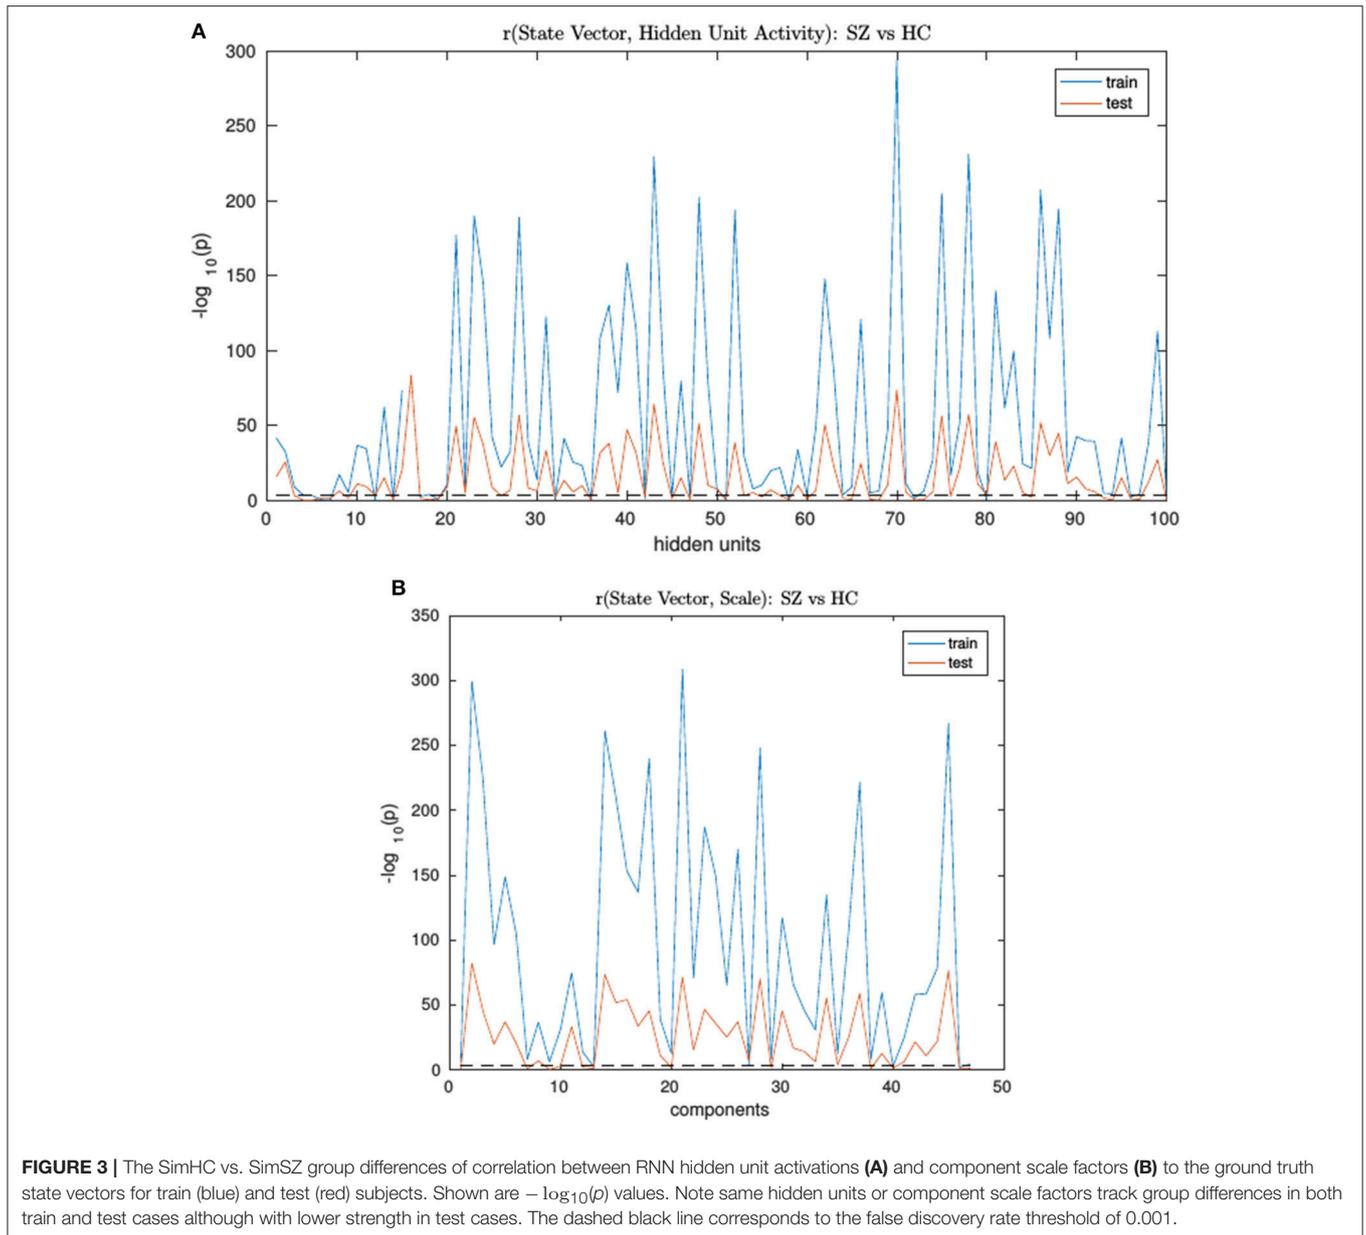
To test the model, we generated synthetic data using SimTB toolbox (Erhardt et al., 2012) in a framework developed to assess dynamics functional connectivity (Allen et al., 2012) and described in Figure 1 of Lehmann et al. (2017). A total of 1,000 subjects corresponding to two groups of subjects, simulated healthy (SimHC) and simulated schizophrenia patients (SimSZ) were generated. A set of 47 time courses were generated for each SimHC and SimSZ subject with the constraint that they have five states (covariance patterns) and a transition probability matrix per group that dictates state transitions derived from data from prior work on real data (Damaraju et al., 2014). The initial state probabilities were also derived from that work. A sequence of 480 time points with a TR of 2 seconds were generated. A total of 1,000 subjects (500 per group) were generated of which first 400 from each group were used during training and remaining 200 samples were used in testing the model. The parameters of hemodynamic response model (delay, undershoot etc) used to simulate the data were also varied per subject to introduce some heterogeneity. The known initial state of a subject and a transition probability matrix that governs transitions ensured a ground

truth state transition vector (a vector of transitions between five simulated states unique to each subject).

An RNN-ICA model was then trained on the 800 subject training data for 500 epochs with model parameters similar to those in subsequent sections. The resultant sources  $\mathbf{S}$ , the source distributions predicted by RNN ( $\mu$ , and  $\sigma$ ), and the RNN hidden unit activations for each subject were then correlated to the subject's ground truth state vector. The trained model was then run on the test data and correlations were again computed between their model outputs and state vectors. We then computed group differences between the correlation distributions of SimHC and SimSZ groups and are summarized for both training and test cases in Figure 3). Our results show that RNN-ICA generalized group differences well to the test set in this setting, as represented in the hidden state activations and scaling factor.

#### 4.2. Task Experiments

To demonstrate the properties and strengths of our model, we apply our method to task fMRI data. Data used in this work is comprised of task-related scans from 28 healthy participants and 24 subjects diagnosed with schizophrenia, all of whom gave written, informed, Hartford hospital and Yale IRB approved consent at the Institute of Living and were compensated for their participation. All participants were scanned during an auditory oddball task (AOD) involving the detection of an infrequent target sound within a series of standard and novel sounds. More



detailed information regarding participant demographics and task details are provided in Swanson et al. (2011).

Scans were acquired at the Olin Neuropsychiatry Research Center at the Institute of Living/Hartford Hospital on a Siemens Allegra 3T dedicated head scanner equipped with 40 mT/m gradients and a standard quadrature head coil. The functional scans were acquired trans-axially using gradient-echo echo-planar-imaging with the following parameters: repeat time (TR) 1.50 s, echo time (TE) 27 ms, field of view 24 cm, acquisition matrix  $[64 \times 64]$ , flip angle  $70^\circ$ , voxel size  $[3.75 \times 3.75 \times 4] \text{mm}^3$ , slice thickness [4]mm, gap [1]mm, 29 slices, ascending acquisition. Six “dummy” scans were acquired at the beginning to allow for longitudinal equilibrium, after which the paradigm was automatically triggered to start by the scanner. The final AOD dataset consisted of 249 volumes for each subject.

Data underwent standard pre-processing steps using the SPM software package (see Calhoun et al., 2008, for further details). Subject scans were masked below a global mean image then each voxel was variance normalized. Each voxel timecourses was then detrended using a 4th-degree polynomial fit, and this was repeated for all subjects. PCA was applied to the complete dataset without whitening, and the first 60 components were kept to reduce the data. Finally, each PCA component had its mean removed before being entered into the model.

#### 4.2.1. Model and Setup

For use in RNNs, the data was then segmented into windowed data, shuffled, and then arranged into random batches. Each PCA loading matrix for subject was comprised of 60 PCA

time courses of length 249. These were segmented into 228 equal-length windowed slices using a window size of 20 and stride of 1. The number of components roughly corresponds to the number found in other studies (Calhoun et al., 2001b, 2008; Allen et al., 2012), and 20 time steps is equivalent to 30 seconds, which has been shown provides a good trade-off in terms of capturing dynamics and not being overly sensitive to noise (Vergara et al., 2017). The final dataset was comprised of 228 volumes for each of the 52 subjects with 60 pca time courses each. These were then randomly shuffled at each epoch into batches of 100 volumes each from random subjects and time points.

We used a simple RNN with 100 recurrent hidden units and a recurrent parameterization as in Equation 7, as we do not anticipate needing to model long range dependencies that necessitate gated models (Hochreiter and Schmidhuber, 1997). The initial hidden state of the RNN was a 2-layer feed forward network with 100 softplus ( $\log(1 + \exp(x))$ ) units using 20% dropout. An additional  $L_2$  decay cost,  $\lambda \sum_{i,j} W_{ij}^2$ , was imposed on the unmixing matrix,  $\mathbf{W}$ , for additional regularization with a decay rate of  $\lambda = 0.002$ . The model was trained using the RMSProp algorithm (Hinton, 2012) with a learning rate of 0.0001 for 500 epochs.

#### 4.2.2. Results

**Figure 4** shows 34 spatial maps back-reconstructed. The spatial maps were filtered from the original 60, omitting white matter, ventricle, and motion artifact features. Each of the spatial maps along with their respective time-courses were sign-flipped to ensure that each back-reconstructed distribution of voxels had positive skew. The maps are highly analogous to those typically found by linear ICA (Calhoun et al., 2001b; Allen et al., 2012), though with more combined positive/negative features in one map.

**Figure 5** shows the functional network connectivity (FNC, Calhoun et al., 2001b; Jafri et al., 2008) matrix, in which the components are grouped according to a multi-level community algorithm Blondel et al. (2008) using the symmetric temporal cross-correlation matrix. For each subject and component, we performed multiple linear regression of the sources,  $\mathbf{s}_{t,n}$ , the predicted means,  $\boldsymbol{\mu}_{t,n}$ , and the predicted scale-factor,  $\boldsymbol{\sigma}_{t,n}$  from each subject to the target and novel stimulus. **Table 1** shows the p-values from a 1-sample  $t$ -test on the beta values across subjects for components with values of  $p \leq 10^{-7}$ . Many components show similar task-relatedness across the source time courses and predicted means, notably temporal gyrus features, parietal lobule, and the default mode network (DMN, which is negatively correlated). In addition, the DMN shows the strongest task-relatedness in the scale factor.

In order to analyze how the RNN encodes dynamics, we analyze the Jacobian of the predicted mean of each component  $i$  at time  $t$  over all components at previous times,  $t'$ :

$$\frac{\partial \mu_{i,t}}{\partial s_{j,t'}} = \sum_k \frac{\partial x_{k,t'}}{\partial s_{j,t'}} \frac{\partial \mu_{i,t}}{\partial x_{k,t'}}. \quad (11)$$

**TABLE 1** |  $p$ -values from a 1-sample  $t$ -test over beta-values with  $p \leq 10^{-7}$  for the target and novel stimulus for the sources,  $\mathbf{s}_{t,n}$ , the predicted means,  $\boldsymbol{\mu}_{t,n}$ , and the predicted scale-factor,  $\boldsymbol{\sigma}_{t,n}$ .

ID	Label	$\boldsymbol{\mu}_i$		$\mathbf{s}_i$		$\boldsymbol{\sigma}_i$
		Targets	Novels	Targets	Novels	Targets
2	FF	3.1e-13 (+)		2.1e-08 (+)		
13	CG	2.2e-13 (+)		1.5e-11 (+)		
14	STG (-)	4.9e-08 (-)		1.4e-08 (-)		
16	Temp.	2.2e-10 (+)		1.9e-17 (+)		
17	L FP (+)	2.6e-11 (+)				
20	Precun. (+)			8.5e-09 (+)		
21	MeFG.	4.1e-08 (+)				
34	DMN	1.0e-18 (-)		1.3e-18 (-)		1.8e-10 (+)
42	IPL	4.7e-19 (+)	4.4e-08 (+)	1.3e-14 (+)	1.1e-08 (+)	
43	MTG (-)	2.1e-19 (+)	7.0e-15 (+)	4.4e-15 (+)	1.2e-14 (+)	
45	AG (+)	3.9e-09 (-)			9.4e-08 (-)	
46	Cere. (+)		5.4e-10 (+)			
47	AG (+)	1.4e-12 (-)	4.5e-10 (+)	6.6e-13 (-)	2.0e-11 (-)	
57	PCing				1.7e-09 (-)	

Beta-values were found for each subject and component using multiple regression to target and novel stimulus, and  $t$ -tests were performed for each component over all subjects. Among the most-significant task-related components to target stimulus include the middle temporal gyrus, default mode network, and the parietal lobule. A legend for ROI label names can be found in the caption of **Figure 4**. The ( $\pm$ ) in the label name specify the sign of the map in **Figure 4**, while the ( $\pm$ ) in the  $p$ -values specifies the sign of the corresponding  $t$ -value.

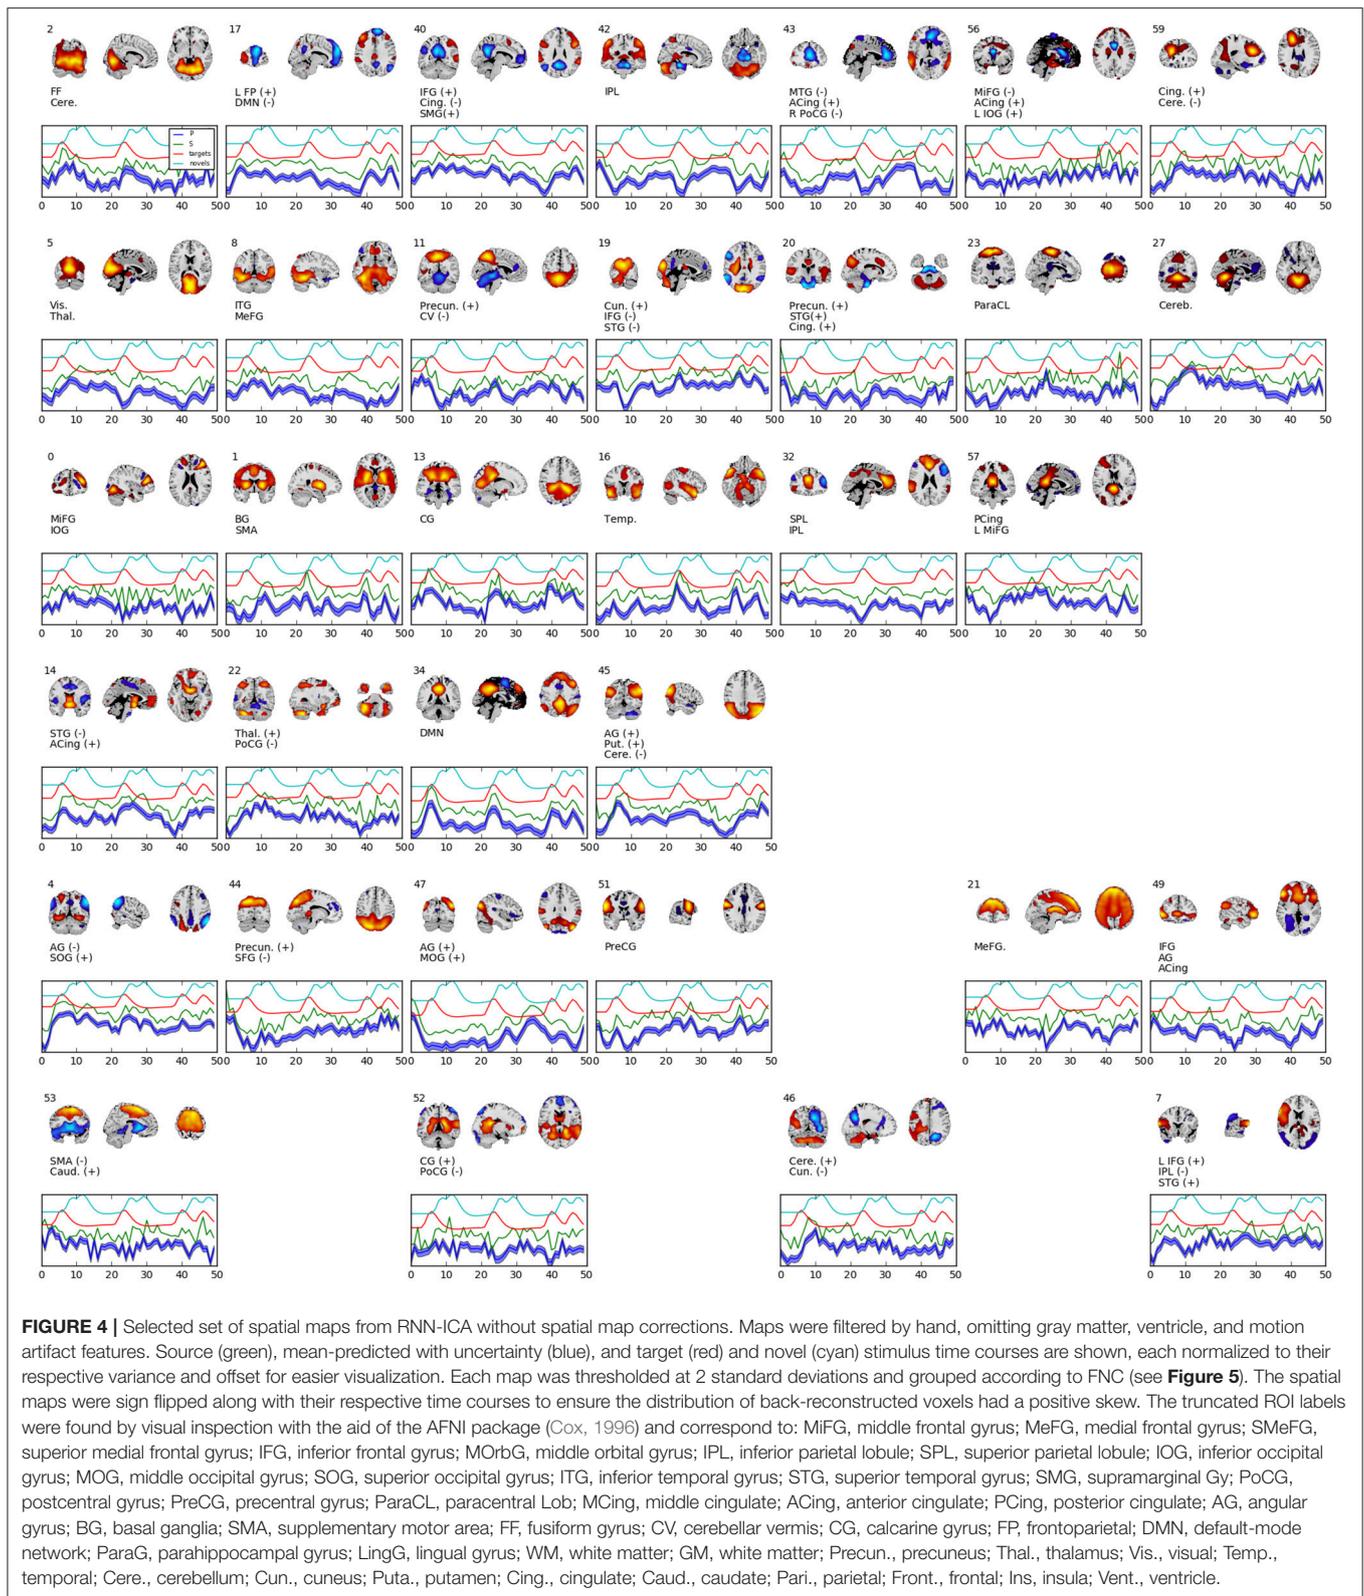
The derivatives are tractable, as the means,  $\mu_{i,t}(\mathbf{x}_{0:t-1})$ , are differentiable functions w.r.t the input  $\mathbf{x}_{0:t-1}$ . These derivatives can be interpreted as being a measure of *directed connectivity* between components in time, as they represent the predicted change of a future component (as understand through the change of its mean value) given change of a previous component. While the full Jacobian provides directed connectivity between source between all pairs of time,  $(t, t')$ , to simplify analysis, we only looked at next-time terms, or  $t' = t - 1$ .

A representative graph is given in **Figure 6**, where the thickness of the edges represents the strength of the directed connection as averaged across time and subjects with the sign removed ( $|\frac{\partial \mu_{i,t}}{\partial s_{j,t'}}|$ ). The color/grouping of the nodes corresponds to the similarity in directed connectivity as measured by the Pearson correlation coefficient:

$$\rho_{i,j} = \frac{\text{Cov}(\bar{v}_i, \bar{v}_j)}{\sigma_{\bar{v}_i} \sigma_{\bar{v}_j}}, \text{ where } \bar{v}_i = \frac{1}{T} \sum_t \frac{\partial \mu_{k,t}}{\partial s_{i,t}}, \quad (12)$$

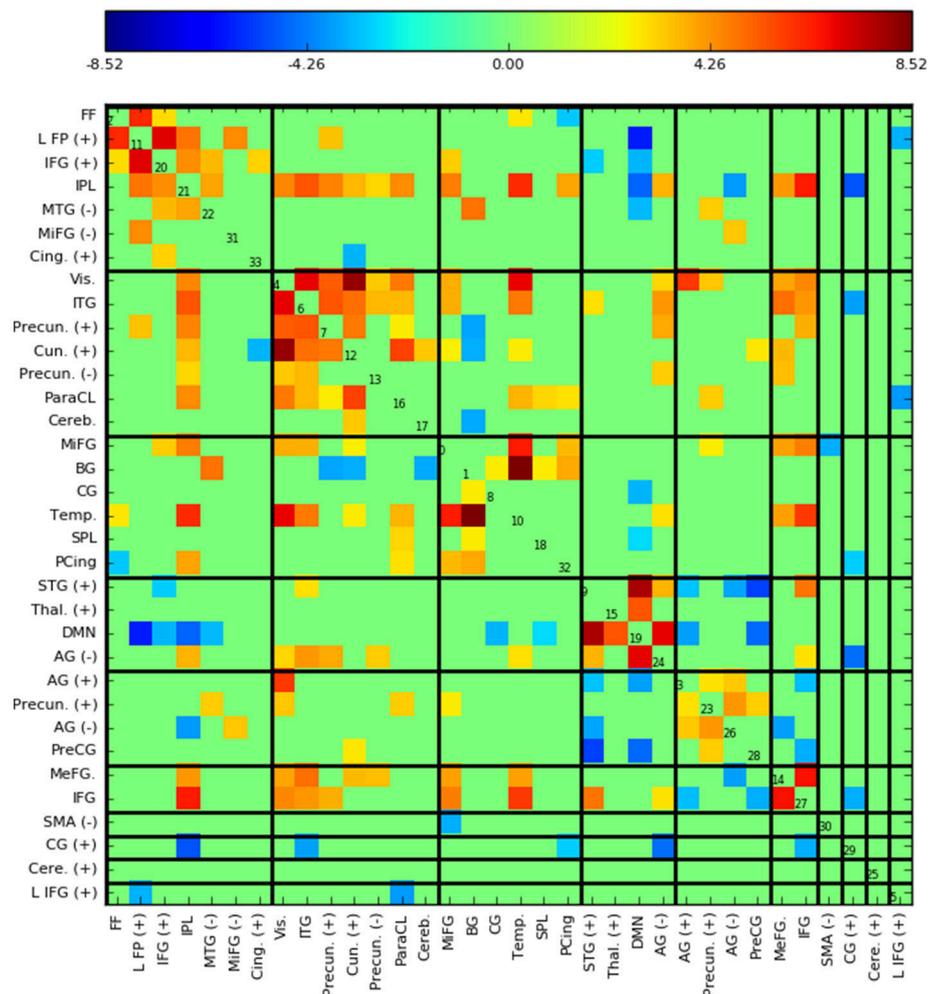
$\text{Cov}(\cdot, \cdot)$  is the covariance, and  $\sigma_{\bar{v}_i}$  is the standard deviation across the components indexed by  $k$ . Grouping was done by constructing an undirected graph using the Pearson coefficients, clustering the vertices using the same community-based hierarchical algorithm as with the FNC above. An example directed connectivity graph with the spatial maps is given in **Figure 7**.

Each of the next-step Jacobian terms were used as time-courses with a multiple-regression to target and novel



stimulus, with significance tested using a one-sample *t*-test as with the time courses and a two-sample *t*-test across groups. The resulting task-related directed connectivity

are represented in **Figure 8** for both targets and novels, with an example graph with spatial maps presented in **Figure 9**. Group-differentiating relationships are given in



**FIGURE 5** | Functional network connectivity (FNC, Jafri et al., 2008) matrix, which is essentially the temporal cross correlation matrix, in this case averaged across subjects. Grouping were found using a multi-level community algorithm Blondel et al. (2008), ordering the FNC axes according the groups. (+) and (-) in the labels indicate regions of interest which have majority positive and negative voxel values, respectively.

Figure 10 with an example graph with spatial maps given in Figure 9.

### 4.3. Resting State Experiments

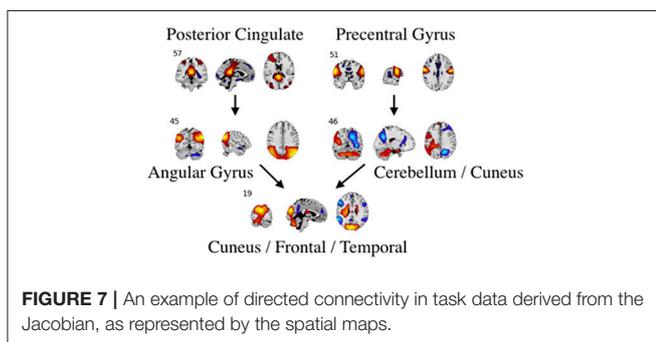
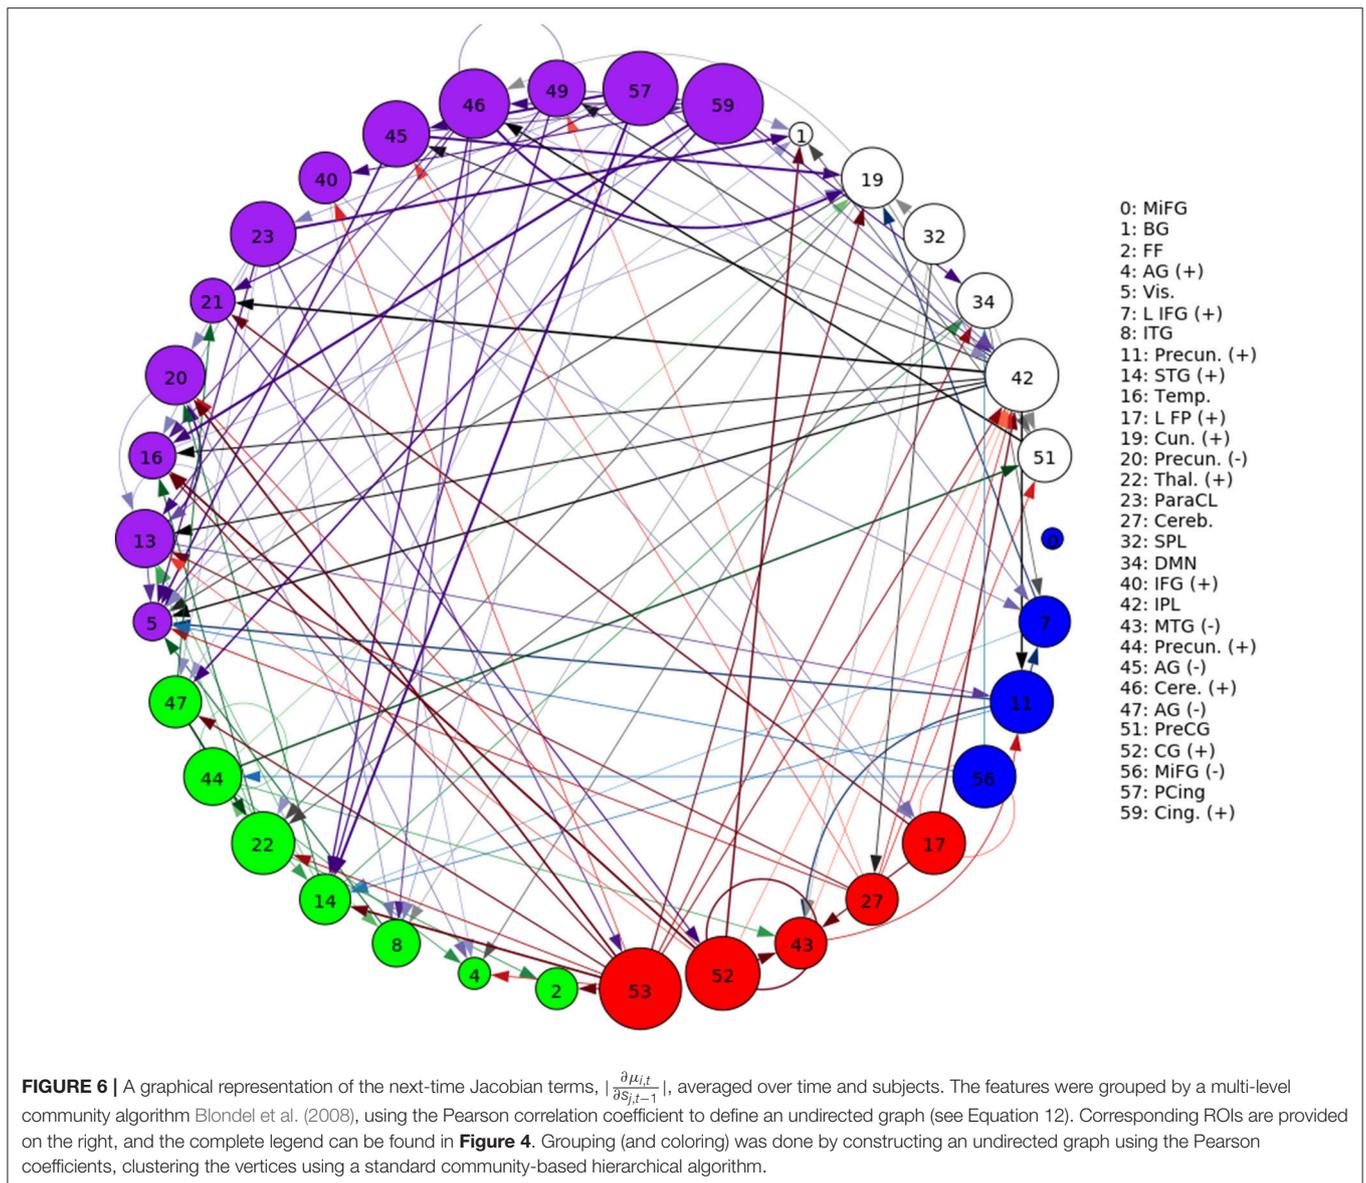
We evaluated our model on resting state data to show RNN-ICA as a viable model and to demonstrate that properties of the network correspond to wake/sleep states. Resting state functional MRI data was collected from 55 subjects for 50 min each (1,505 volumes, TR= 2.08 s) with a Siemens 3T Trio scanner while the subjects transitioned from wakefulness to at most sleep stage N3 (see Tagliazucchi et al., 2012, for more details). This data was approved by ethics committee of Goethe University. Simultaneous EEG was acquired facilitating sleep staging per AASM criteria resulting in a hypnogram per subject (a vector assignment of consecutive 30 s EEG epochs to one of wakeful(W), N1, N2 and N3 sleep stae). We discarded first 5 time points to account for T1 equilibration effects.

After performing rigid body realignment and slice-timing correction, subject data was warped to MNI space using SPM12.

Then voxel time courses were despiked using AFNI. We then regressed out voxel time courses with respect to their head motion parameters (and their derivatives and squares), their mean white matter and CSF signals. Next, we bandpass filtered the data with a passband of 0.01–0.15 Hz. We extracted mean ROI time courses from 268 nodes extracted from the bioimage suite (Papademetris et al., 2006) and reported in Shen et al. (2013).

#### 4.3.1. Model and Setup

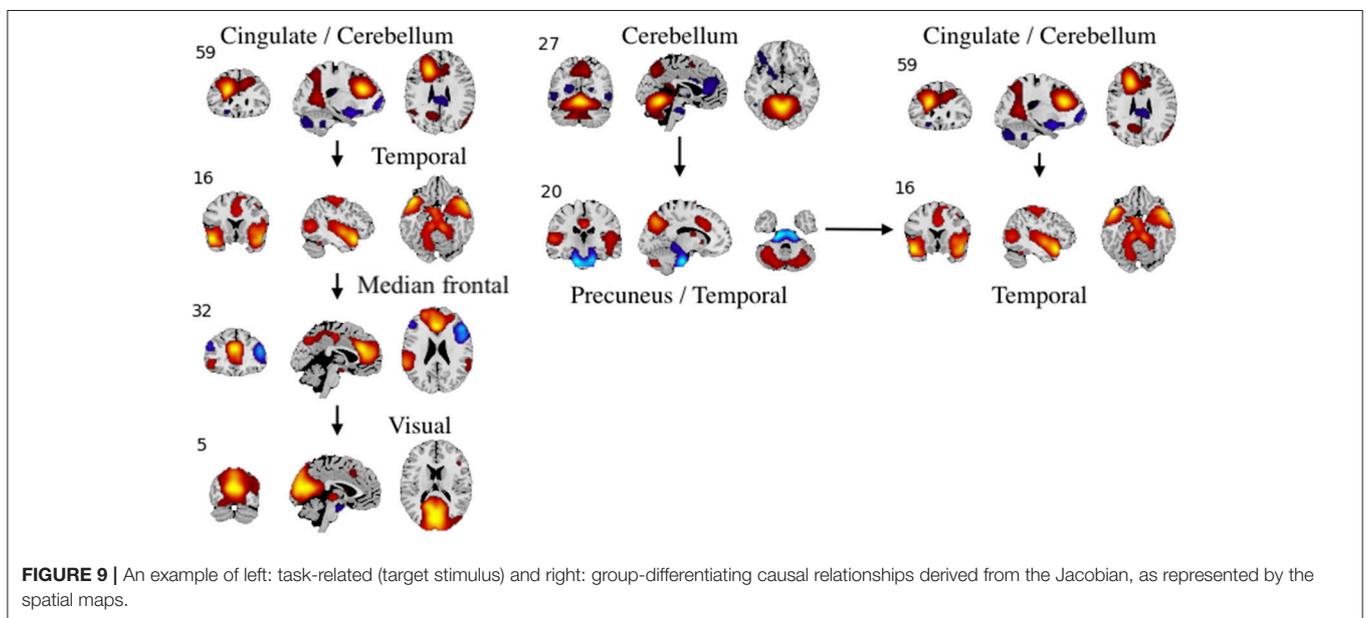
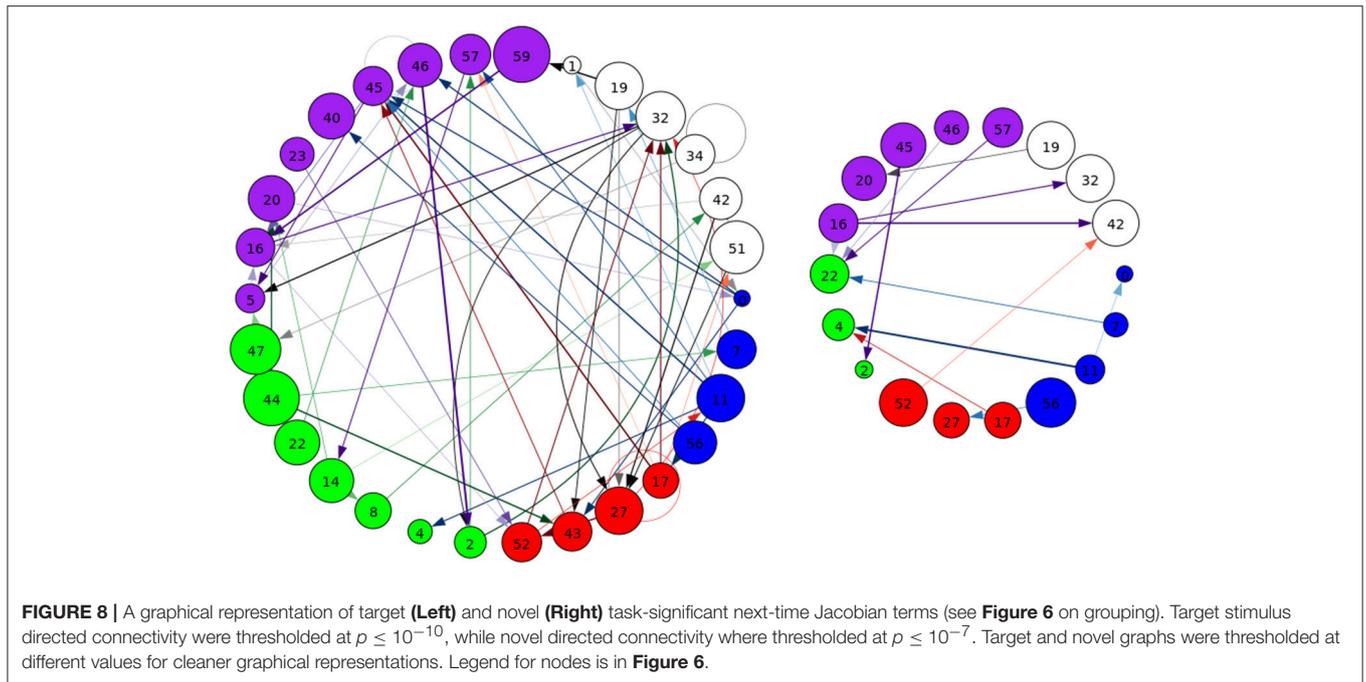
We used the same model and training procedure as with our task data analysis in the previous section. Of the 55 subjects, 50 subjects were used during training and 5 subjects were left out for testing. We then examined the correspondence between hidden recurrent units of the trained model and subject hypnogram as well as between mean and scale of predictive source distribution and hypnogram. Similar tests were run on the model outputs on the 5 left out test cases.



**4.3.2. Results**

The activity of several hidden recurrent units of trained model was predictive of wakefulness across all subjects (see **Figure 11**

for an example subject). The RNN hidden unit activity (bound between -1 and 1) stays at the extremes during awake state exhibiting higher standard deviation and the activity tends toward zero with lower standard deviation as the subject transitions from wakefulness to sleep. One-way ANOVA on the absolute mean and standard deviation of hidden unit activity by hypnogram state shows significant group differences in mean ( $p \leq 10^{-29}$ ) and standard deviation ( $p \leq 10^{-14}$ ). Subsequent *post-hoc t*-tests reveal significant reductions in both from wakefulness and light sleep N1 state to deeper sleep stages N2 and N3 states, and also between N2 and N3 states (means:[0.6642 0.6554 0.4558 0.2033], and standard deviations: [0.1868 0.1997 0.1567 0.0579]; all these p-values  $\leq 10^{-5}$  after correcting for multiple comparisons). In addition, the scaling factor tended to correlate well with changes of state, as measures by correlation with a smoothed derivative of the



hypnogram. Figure 12 shows the correlation coefficients between RNN hidden units to subject hypnogram state, component scale factors,  $\sigma$  to subject hypnogram vector. Several hidden states show consistent correlation to hypnograms, indicating the RNN is encoding subject sleep state. Similarly some component scale factors also encode sleep states. Surprisingly, however, the source time courses,  $\mathbf{s}$ , and the means,  $\mu$ , did not. Finally, some component scale factors correlate somewhat consistently with changes in state across subjects. This indicates that the model is encoding changes of state in terms of uncertainty.

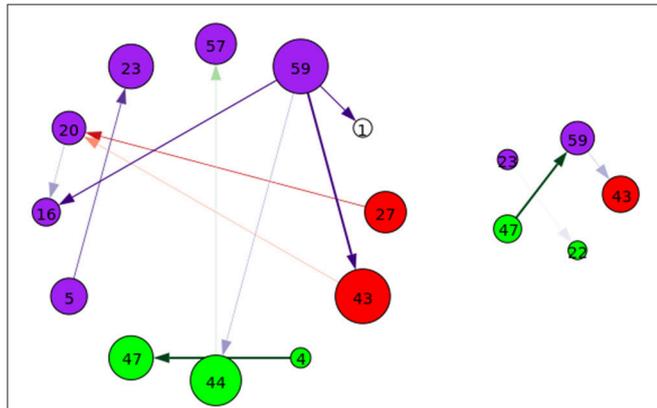
## 5. DISCUSSION AND CONCLUSION

### 5.1. Summary

In this work, we demonstrate how RNNs can be used to separate conditionally independent sources analogous to independent component analysis but with the benefits of modeling temporal dynamics through recurrent parameters. Results show that this approach is effective for modeling both task-related and resting-state functional magnetic imaging (fMRI) data. Using this approach, we are able to separate similar components to ICA,

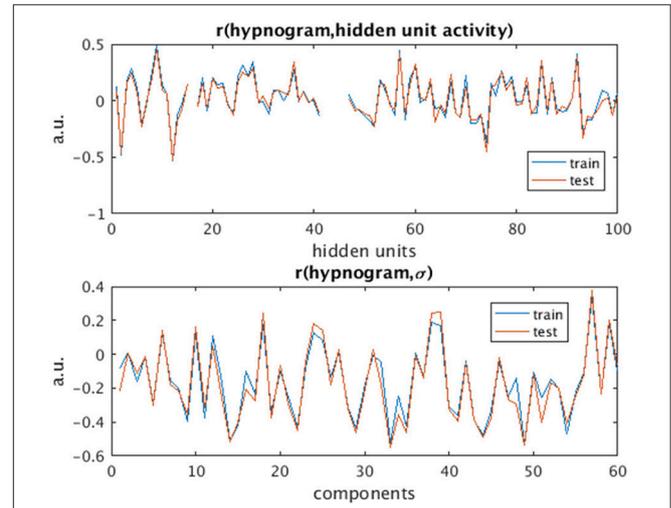
but having the additional benefit of directly analyzing temporal dynamics through the recurrent parameters.

Notably, in addition to finding similar maps and task-relatedness as with ICA, we are able to derive directed temporal connectivity which is task-related and group-differential, and these are derived directly from the parameters of the RNN.

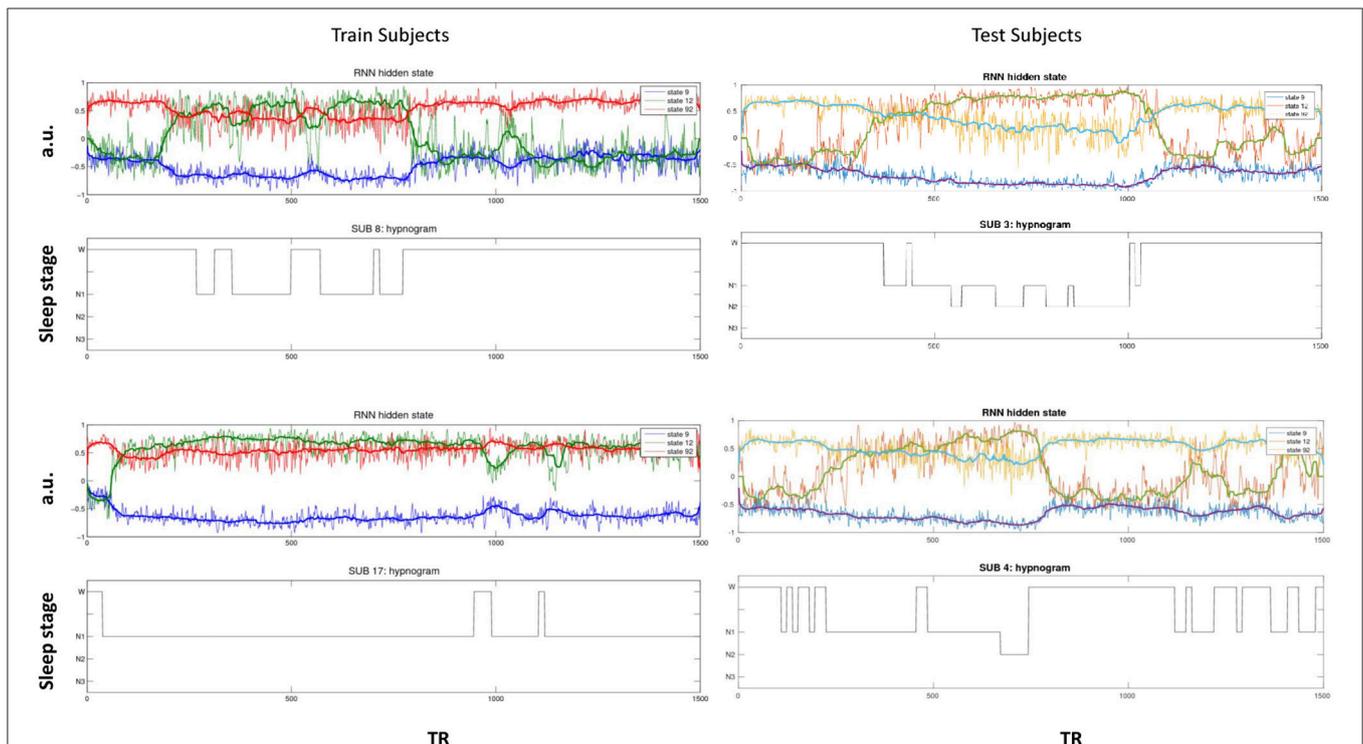


**FIGURE 10 |** A graphical representation of target (Left) and novel (Right) group-differentiating next-time Jacobian terms (see Figure 6 on grouping). Target stimulus directed connectivity were thresholded at  $p \leq 0.001$ . Legend for nodes is in Figure 6. This shows that the influence between components across time when different stimulus is present can vary across groups.

In addition, for resting state data, we found that some hidden unit activity corresponded very well with wake/sleep states and that the uncertainty factor was consistent with



**FIGURE 12 |** Correlation values of subject hypnogram to learned RNN model hidden units (Top) and to component scale factors  $\sigma$  (Bottom) for train (blue) and test (red) subjects. The same units and scale factors track subject neurobiological state as derived from EEG for both train and test cases. Note some hidden unit activations are flat and so the correlation value is empty.



**FIGURE 11 |** Select hidden unit time courses with the corresponding hypnogram for two subjects in the training data (Left) and two test subjects (Right) not used while training. Hidden states track subject neurobiological state [Wakeful (W) or N1, N2, and N3 sleep stages] just using fMRI activity. The bold lines are the median filtered activation time courses.

changes of state, both of which were learned in a completely unsupervised way.

## 5.2. Related Work

Our method introduces deep and non-linear computations in time to MLE independent component analysis (MLE ICA) without sacrificing the simplicity of linear relationships between source and observation. MLE ICA has an equivalent learning objective to infomax ICA, widely used in fMRI studies, in which the sources are drawn from a factorized logistic distribution (Hyvärinen et al., 2004). While the model learns a linear transformation between data and sources through the unmixing matrix, the source dynamics are encoded by a deep non-linear transformation with recurrent structure, as represented by an RNN. Alternative non-linear parameterizations of the ICA transformation exist that use deep neural networks have been shown to work with fMRI data (Castro et al., 2016). Such approaches allow for deep and non-linear static spatial maps and are compatible with our learning objective. Temporal ICA as used in group ICA (Calhoun et al., 2009), like spatial ICA, does capture some temporal dynamics, but only as summaries through a one- to two-stage PCA preprocessing step. These temporal summaries are captured and can be analyzed, however they are not learned as part of an end-to-end learning objective. Overall, the strengths of RNN-ICA compared to these methods are the dynamics are directly learned as model parameters, which allows for richer and higher-order temporal analyses, as we showed in the previous section.

Recurrent neural networks do not typically incorporate latent variables, as this requires expensive inference. Versions that incorporate stochastic latent variables exist, are trainable via variational methods, and working approaches for sequential data exist (Chung et al., 2015). However, these require complex inference which introduces variance into learning that may make training with fMRI data challenging. Our method instead incorporates concepts from noiseless ICA, which reduces

inference to the inverse of a generative transformation. The consequence is that the temporal analyses are relatively simple, relying on only the tractable computation of the Jacobian of component conditional densities given the activations.

## 5.3. Future Work

The RNN-ICA model provides a unique mode of analysis previously unavailable to fMRI research. Results are encouraging, in that we were able to find both task-related and group-differentiating directed connectivity, however the broader potential of this approach is unexplored. It is our belief that this method will expand neuroscience research that involves temporal data, leading to new and significant conclusions.

Finally, the uncertainty factor in our resting state experiments may indicate a novel application for imaging data through RNN-ICA, that is change-of-state detection. The model we employed was simple, as was not intended to take advantage of this. It is quite possible that further modifications could produce a model that reliably predicts change-of-state in fMRI and EEG data.

## AUTHOR CONTRIBUTIONS

RH performed contributed the primary research idea, implementation, experimentation, organization, and primary writing. ED performed experiments on the resting state MRI as well as wrote the experiment and results sections there. KC provided insights related to RNNs. HL provided feedback related to the method and his work on resting state MRI. SP helped with writing as well as research direction. VC was the PI and provided funding, feedback, helped with writing, and research direction.

## FUNDING

This work was supported in part by National Institutes of Health grants 2R01EB005846, P20GM103472, and R01EB020407 and National Science Foundation grant #1539067.

## REFERENCES

- Allen, E. A., Erhardt, E. B., Wei, Y., Eichele, T., and Calhoun, V. D. (2012). Capturing inter-subject variability with group independent component analysis of fMRI data: a simulation study. *Neuroimage* 59, 4141–4159. doi: 10.1016/j.neuroimage.2011.10.010
- Bahdanau, D., Cho, K., and Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Bell, A. J., and Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Comput.* 7, 1129–1159.
- Biswal, B., Zerrin Yetkin, F., Haughton, V. M., and Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* 34, 537–541.
- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *J. Stat. Mech. Theor. Exp.* 2008:P10008. doi: 10.1088/1742-5468/2008/10/P10008
- Calhoun, V. D., and Adali, T. (2012). Multisubject independent component analysis of fmri: a decade of intrinsic networks, default mode, and neurodiagnostic discovery. *IEEE Rev. Biomed. Eng.* 5, 60–73. doi: 10.1109/RBME.2012.2211076
- Calhoun, V. D., Adali, T., Pearlson, G., and Pekar, J. (2001a). Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. *Hum. Brain Mapp.* 13, 43–53. doi: 10.1002/hbm.1024
- Calhoun, V. D., Adali, T., Pearlson, G. D., and Pekar, J. J. (2001b). A method for making group inferences from functional MRI data using independent component analysis. *Hum. Brain Mapp.* 14, 140–151. doi: 10.1002/hbm.1048
- Calhoun, V. D., Kiehl, K. A., and Pearlson, G. D. (2008). Modulation of temporally coherent brain networks estimated using ICA at rest and during cognitive tasks. *Human Brain Mapping*, 29, 828–838. doi: 10.1002/hbm.20581
- Calhoun, V. D., Liu, J., and Adali, T. (2009). A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data. *Neuroimage* 45, S163–S172. doi: 10.1016/j.neuroimage.2008.10.057
- Castro, E., Hjelm, R. D., Plis, S., Dihn, L., Turner, J., and Calhoun, V. (2016). Deep independence network analysis of structural brain imaging: application to schizophrenia. *IEEE Trans. Med. Imaging*. 35, 1729–1740. doi: 10.1109/TMI.2016.2527717
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.

- Chung, J., Kastner, K., Dinh, L., Goel, K., Courville, A. C., and Bengio, Y. (2015). "A recurrent latent variable model for sequential data," in *Advances in neural information processing systems*, 2980–2988.
- Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res.* 29, 162–173.
- Damaraju, E., Allen, E. A., Belger, A., Ford, J., McEwen, S., Mathalon, D., et al. (2014). Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. *Neuroimage Clin.* 5, 298–308. doi: 10.1016/j.nicl.2014.07.003
- Damoiseaux, J., Rombouts, S., Barkhof, F., Scheltens, P., Stam, C., Smith, S. M., et al. (2006). Consistent resting-state networks across healthy subjects. *Proc. Natl. Acad. Sci. U.S.A.* 103, 13848–13853. doi: 10.1073/pnas.0601417103
- Doucet, A., De Freitas, N., and Gordon, N. (eds.). (2001). "An introduction to sequential Monte Carlo methods," in *Sequential Monte Carlo Methods in Practice* (New York, NY: Springer Science+Business Media), 3–14.
- Erhardt, E. B., Allen, E. A., Wei, Y., Eichele, T., and Calhoun, V. D. (2012). SimTB, a simulation toolbox for fMRI data under a model of spatiotemporal separability. *Neuroimage* 59, 4160–4167. doi: 10.1016/j.neuroimage.2011.11.088
- Fu, G.-S., Phlypo, R., Anderson, M., and Adali, T. (2014). Blind source separation by entropy rate minimization. *IEEE Trans. Signal Process.* 62, 4245–4255. doi: 10.1109/TSP.2014.2333563
- Graves, A. (2013). Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.
- Güçlü, U. and van Gerven, M. A. (2017). Modeling the dynamics of human brain activity with recurrent neural networks. *Front. Comput. Neurosci.* 11:7. doi: 10.3389/fncom.2017.00007
- Hinton, G. (2012). *Neural Networks for Machine Learning*. Coursera. Available online at: <https://www.coursera.org/learn/neural-networks>
- Hjelm, R. D., Calhoun, V. D., Salakhutdinov, R., Allen, E. A., Adali, T., and Plis, S. M. (2014). Restricted boltzmann machines for neuroimaging: an application in identifying intrinsic networks. *Neuroimage* 96, 245–260. doi: 10.1016/j.neuroimage.2014.03.048
- Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9, 1735–1780.
- Hyvärinen, A., Karhunen, J., and Oja, E. (2004). *Independent Component Analysis*, Vol. 46. New York, NY; Chichester; Weinheim; Brisbane; Singapore; Toronto: John Wiley & Sons.
- Hyvärinen, A. and Oja, E. (2000). Independent component analysis: algorithms and applications. *Neural Netw.* 13, 411–430. doi: 10.1016/S0893-6080(00)00026-5
- Jafri, M. J., Pearlson, G. D., Stevens, M., and Calhoun, V. D. (2008). A method for functional network connectivity among spatially independent resting-state components in schizophrenia. *Neuroimage* 39, 1666–1681. doi: 10.1016/j.neuroimage.2007.11.001
- Kim, D., Burge, J., Lane, T., Pearlson, G. D., Kiehl, K. A., and Calhoun, V. D. (2008). Hybrid ICA–bayesian network approach reveals distinct effective connectivity differences in schizophrenia. *Neuroimage* 42, 1560–1568. doi: 10.1016/j.neuroimage.2008.05.065
- Lehmann, B., White, S., Henson, R., Cam-CAN, and Geerligns, L. (2017). Assessing dynamic functional connectivity in heterogeneous samples. *Neuroimage* 157, 635–647. doi: 10.1016/j.neuroimage.2017.05.065
- Papademetris, X., Jackowski, M. P., Rajeevan, N., DiStasio, M., Okuda, H., Constable, R. T., et al. (2006). Bioimage suite: an integrated medical image analysis suite: an update. *Insight J.* 2006:209.
- Plis, S. M., Hjelm, R. D., Salakhutdinov, R., and Calhoun, V. D. (2014). Deep learning for neuroimaging: a validation study. *Front. Neurosci.* 8:229. doi: 10.3389/fnins.2014.00229
- Shen, X., Tokoglu, F., Papademetris, X., and Constable, R. T. (2013). Groupwise whole-brain parcellation from resting-state fmri data for network node identification. *Neuroimage* 82, 403–415. doi: 10.1016/j.neuroimage.2013.05.081
- Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., et al. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci. U.S.A.* 106, 13040–13045. doi: 10.1073/pnas.0905267106
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). "Sequence to sequence learning with neural networks," in *Advances in Neural Information Processing Systems*, 3104–3112.
- Swanson, N., Eichele, T., Pearlson, G., Kiehl, K., Yu, Q., and Calhoun, V. D. (2011). Lateral differences in the default mode network in healthy controls and patients with schizophrenia. *Hum. Brain Mapp.* 32, 654–664. doi: 10.1002/hbm.21055
- Tagliazucchi, E., von Wegner, F., Morzelewski, A., Borisov, S., Jahnke, K., and Laufs, H. (2012). Automatic sleep staging using fMRI functional connectivity data. *Neuroimage* 63, 63–72.
- Van Den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., et al. (2016). Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*.
- Vergara, V. M., Mayer, A. R., Damaraju, E., and Calhoun, V. D. (2017). The effect of preprocessing in dynamic functional network connectivity used to classify mild traumatic brain injury. *Brain Behav.* 7:e00809. doi: 10.1002/brb.3.809
- Zuo, X.-N., Kelly, C., Adelman, J. S., Klein, D. F., Castellanos, F. X., and Milham, M. P. (2010). Reliable intrinsic connectivity networks: test–retest evaluation using ICA and dual regression approach. *Neuroimage* 49, 2163–2177. doi: 10.1016/j.neuroimage.2009.10.080

**Conflict of Interest Statement:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2018 Hjelm, Damaraju, Cho, Laufs, Plis and Calhoun. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.