Generating Hypotheses of Dynamic Causal Graphs in Neuroscience: Leveraging Generative Factor Models of Observed Time Series

Zachary C. Brown (1), David Carlson (1,2)

(1) Department of Electrical & Computer Engineering *

(2) Departments of Civil & Environmental Engineering; Biostatistics & Bioinformatics; & Computer Science *

*Duke University, Durham, NC 15213, USA.

Motivation and Assumptions

We seek to reduce the number of hypotheses that must be tested before finding causal relationships in neuroscientific applications.

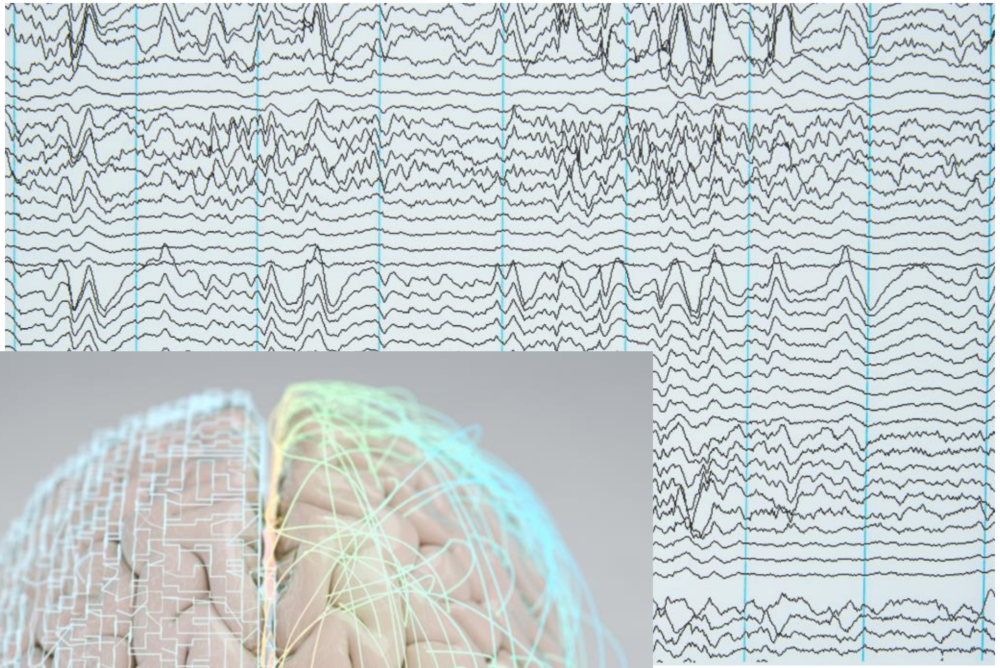

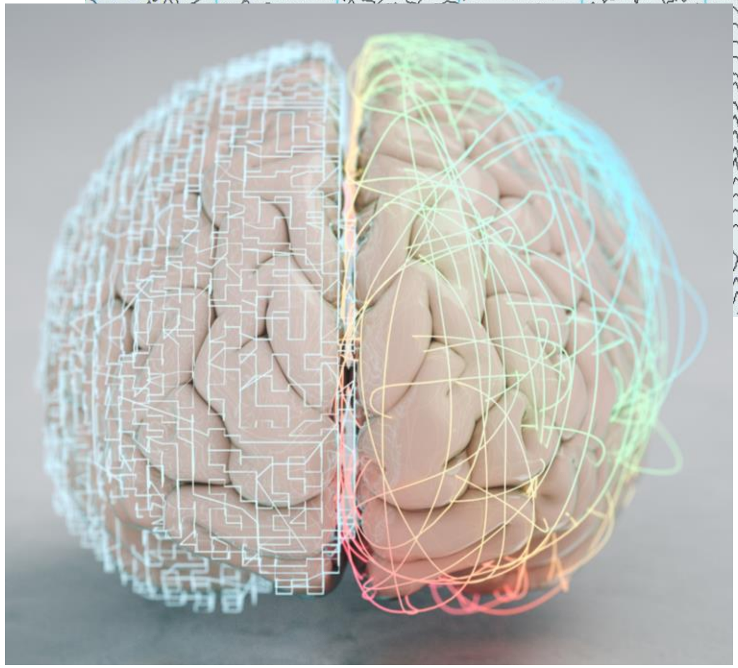

Motivation and Assumptions

We seek to reduce the number of hypotheses that must be tested before finding causal relationships in neuroscientific applications. As such, we present :

1. a novel approach to formulating hypotheses using deep generative factor models
2. methods for including auxiliary (behavioral) labels

We only assume data consists of regularly sampled (noisy) time series. We make no assumptions about underlying generative processes

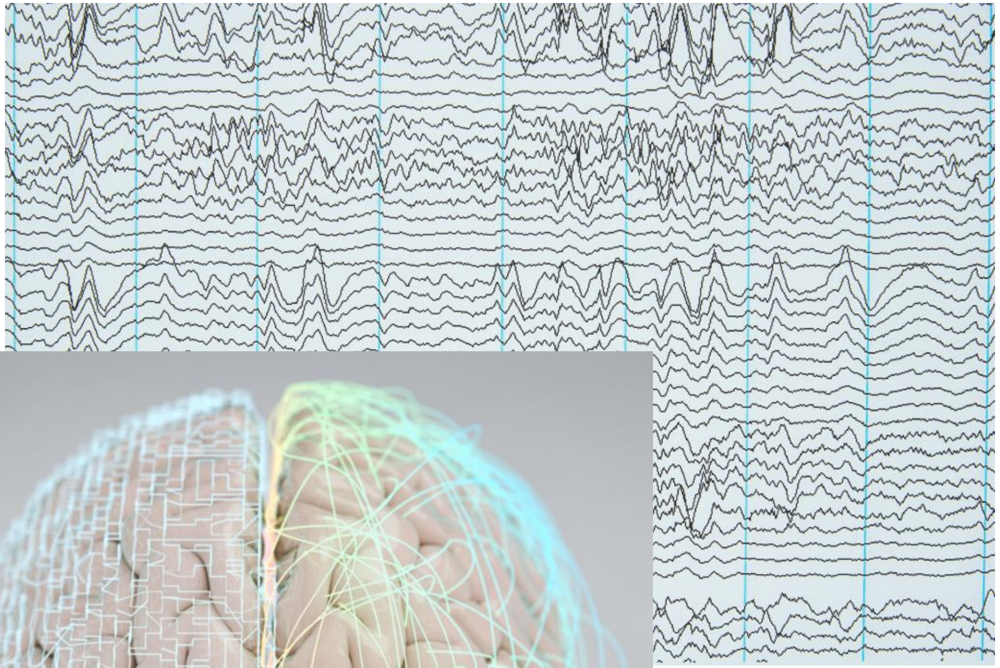

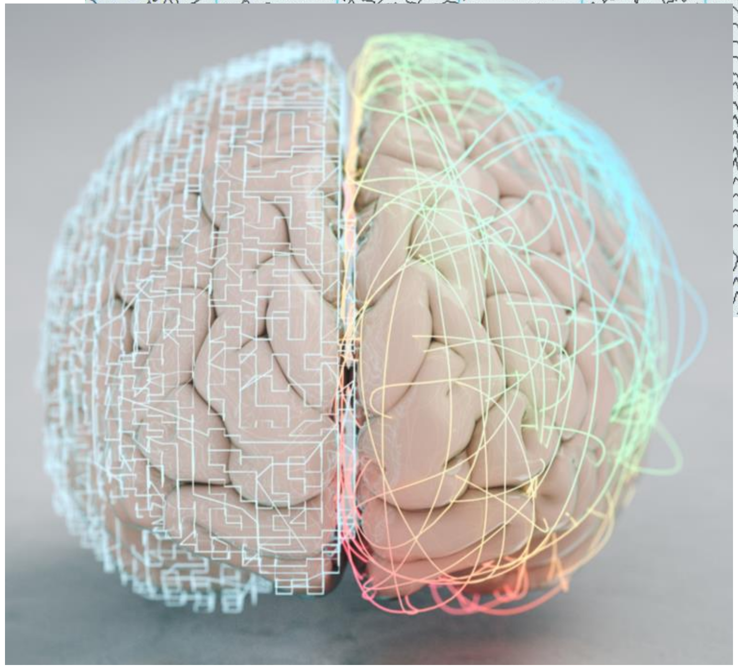

Methods: REDCLIFF-S Overview

REDCLIFF-S divides hypothesis generation into two tasks: individual factor learning and factor relevance prediction.

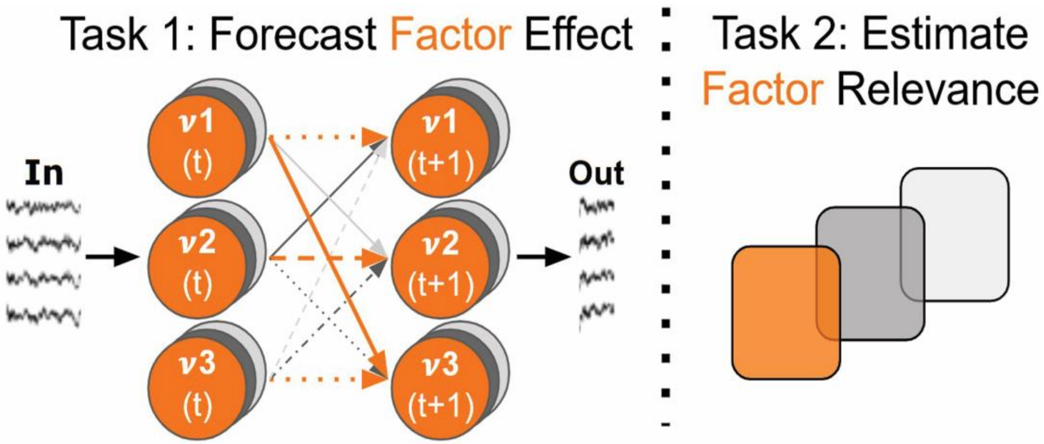

Methods: REDCLIFF-S Overview

REDCLIFF-S divides hypothesis generation into two tasks: individual factor learning and factor relevance prediction.

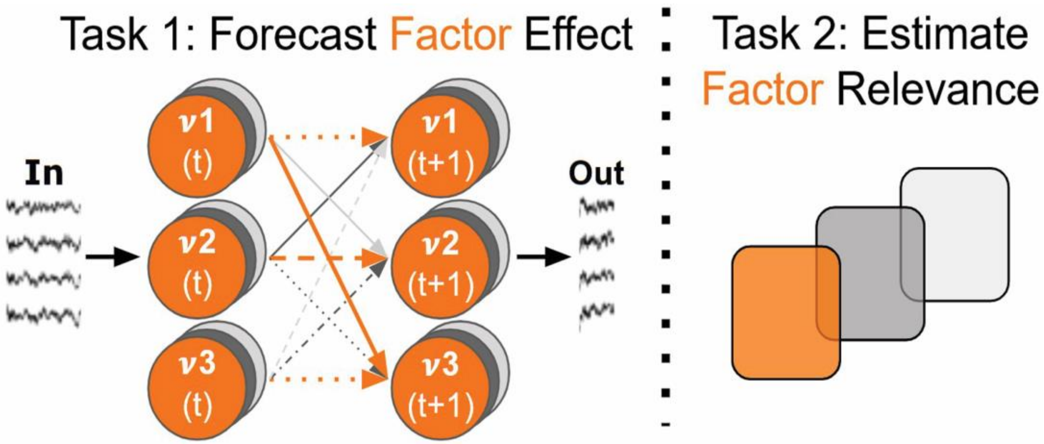

This distinction between tasks is reflected in the separation between generative factors and the state model.

$$\hat{\mathbf{x}}_{n,t+1} = \sum_{k=1}^{\mathrm{n_k}} \alpha_{n,k,t+1} f_{\phi_k}(\mathbf{x}_{n,t-\tau_{\mathrm{in}}:t})$$

$$g_\theta(\mathbf{x}) = \begin{bmatrix} \boldsymbol{\alpha} \\ \hat{\mathbf{y}} \end{bmatrix}$$

Methods: REDCLIFF-S Structure

Both generative factors and the state model have direct access to the ‘input’ data. This allows us to (linearly) identify how each factor depends on system history even if its generative role (or ‘relevance’) is nonlinear in nature.

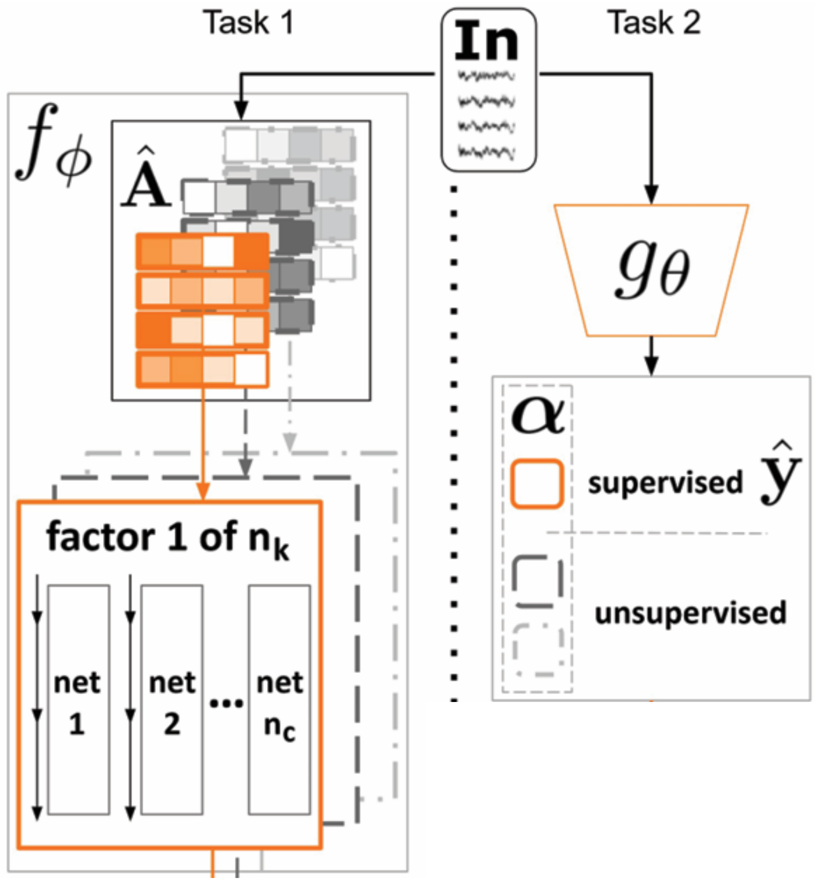

Methods: REDCLIFF-S Structure

Both generative factors and the state model have direct access to the 'input' data. This allows us to (linearly) identify how each factor depends on system history even if its generative role (or 'relevance') is nonlinear in nature.

We forecast the system's evolution with a weighted superposition of each factor's forecasted effect based on the predicted system state

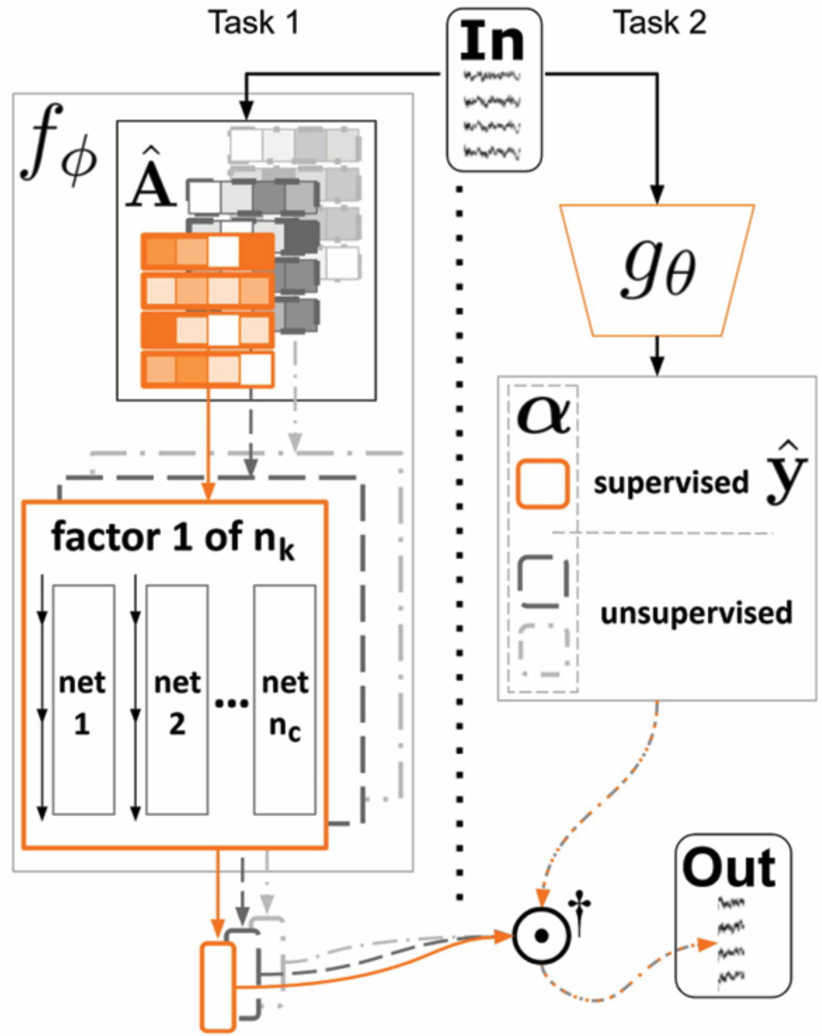

Methods: REDCLIFF-S Objective Function

System Evolution Forecasting Term

$$\boxed{\omega \mathrm{MSE}(\mathbf{X}, \hat{\mathbf{X}})}$$

Methods: REDCLIFF-S Objective Function

Penalty for Temporal Length of Predicted Causal Relationships

$$\omega \mathrm{MSE}(\mathbf{X}, \hat{\mathbf{X}}) + \boxed{\eta \sum_{k=1}^{\mathrm{n_k}} \sum_{t=1}^{\tau_{\mathrm{in}}} \log(t+1) ||{}^{k}\hat{\mathbf{A}}_{:,:,t}||_1}$$

Methods: REDCLIFF-S Objective Function

$$\omega \mathrm{MSE}(\mathbf{X}, \hat{\mathbf{X}}) + \eta \sum_{k=1}^{\mathrm{n_k}} \sum_{t=1}^{\tau_{\mathrm{in}}} \log(t+1) ||^{k}\hat{\mathbf{A}}_{:,:,t}||_1$$

$$+ \boxed{\rho \sum_{p=1}^{\mathrm{n_k}} \sum_{q=(p+1)}^{\mathrm{n_k}} \mathrm{CosSim}(^{p}\tilde{\mathbf{A}} - \mathbf{I}, ^{q}\tilde{\mathbf{A}} - \mathbf{I})}$$

Factor Causal Prediction Similarity Penalty

Methods: REDCLIFF-S Objective Function

$$\mathcal{L}_f(\phi, \mathcal{D}) = \omega \mathrm{MSE}(\mathbf{X}, \hat{\mathbf{X}}) + \eta \sum_{k=1}^{\mathrm{n_k}} \sum_{t=1}^{\tau_{\mathrm{in}}} \log(t+1) ||^k \hat{\mathbf{A}}_{:,:,t}||_1$$

$$+ \rho \sum_{p=1}^{\mathrm{n_k}} \sum_{q=(p+1)}^{\mathrm{n_k}} \mathrm{CosSim}(^p \tilde{\mathbf{A}} - \mathbf{I}, ^q \tilde{\mathbf{A}} - \mathbf{I})$$

Methods: REDCLIFF-S Objective Function

$$\mathcal{L}_f(\phi, \mathcal{D}) = \omega \text{MSE}(\mathbf{X}, \hat{\mathbf{X}}) + \eta \sum_{k=1}^{\mathrm{n_k}} \sum_{t=1}^{\tau_{\mathrm{in}}} \log(t+1) ||^k \hat{\mathbf{A}}_{:,:,t}||_1$$

$$+ \rho \sum_{p=1}^{\mathrm{n_k}} \sum_{q=(p+1)}^{\mathrm{n_k}} \text{CosSim}(^p\tilde{\mathbf{A}} - \mathbf{I}, ^q\tilde{\mathbf{A}} - \mathbf{I})$$

$$\mathcal{L}(\phi, \theta, \mathcal{D}) = \mathcal{L}_f(\phi, \mathcal{D}) + \boxed{\gamma\Big(-1 + \sum_{n=1}^{N} ||\boldsymbol{\alpha}_{n,:}||_1\Big)} + \boxed{\lambda \text{MSE}(\mathbf{Y}, \hat{\mathbf{Y}})}$$

Factor Weighting Sparsity Penalty

System State Prediction Term

Experiments on Synthetic Data: Highlights

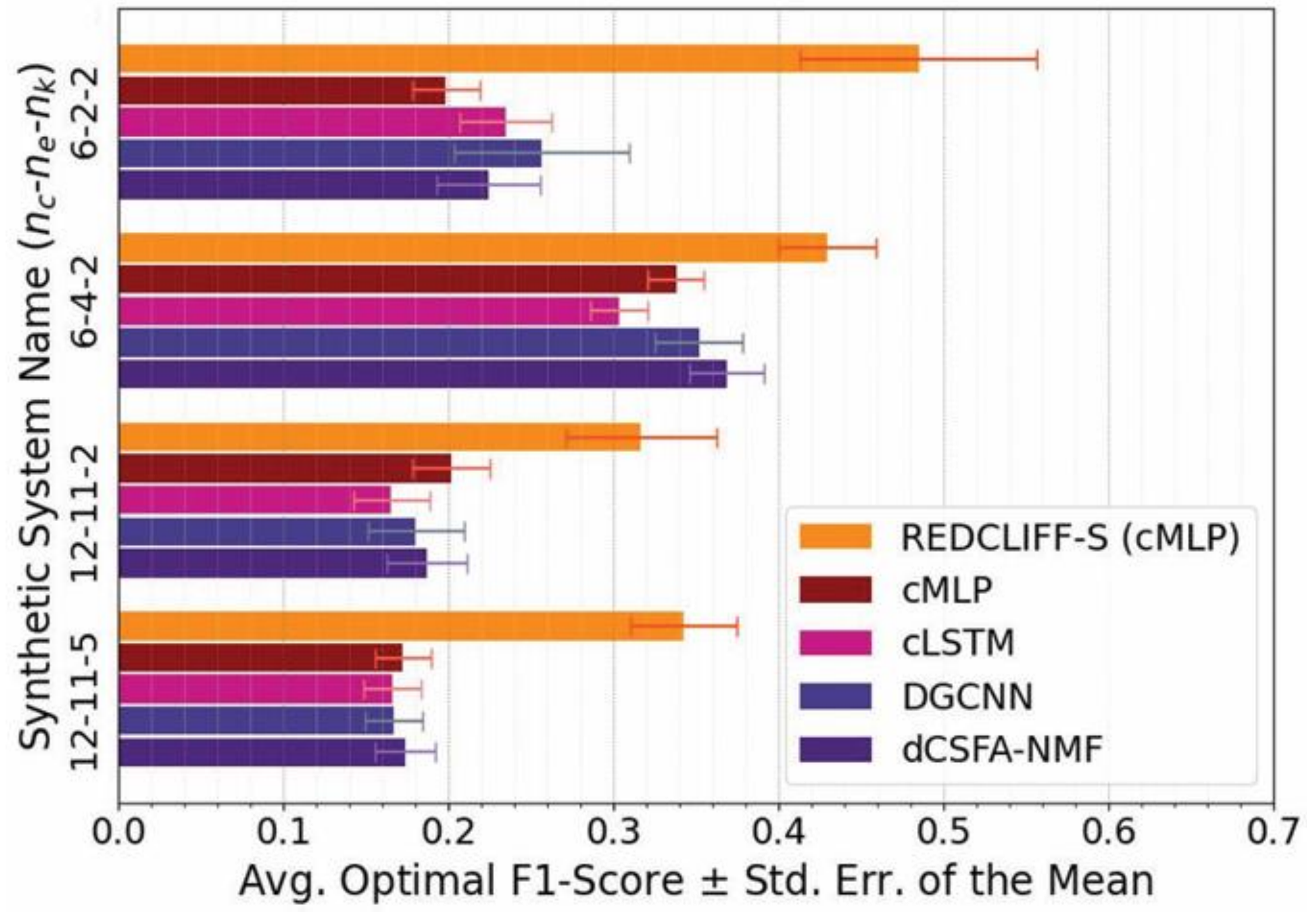

Experiments on Synthetic Data: Highlights

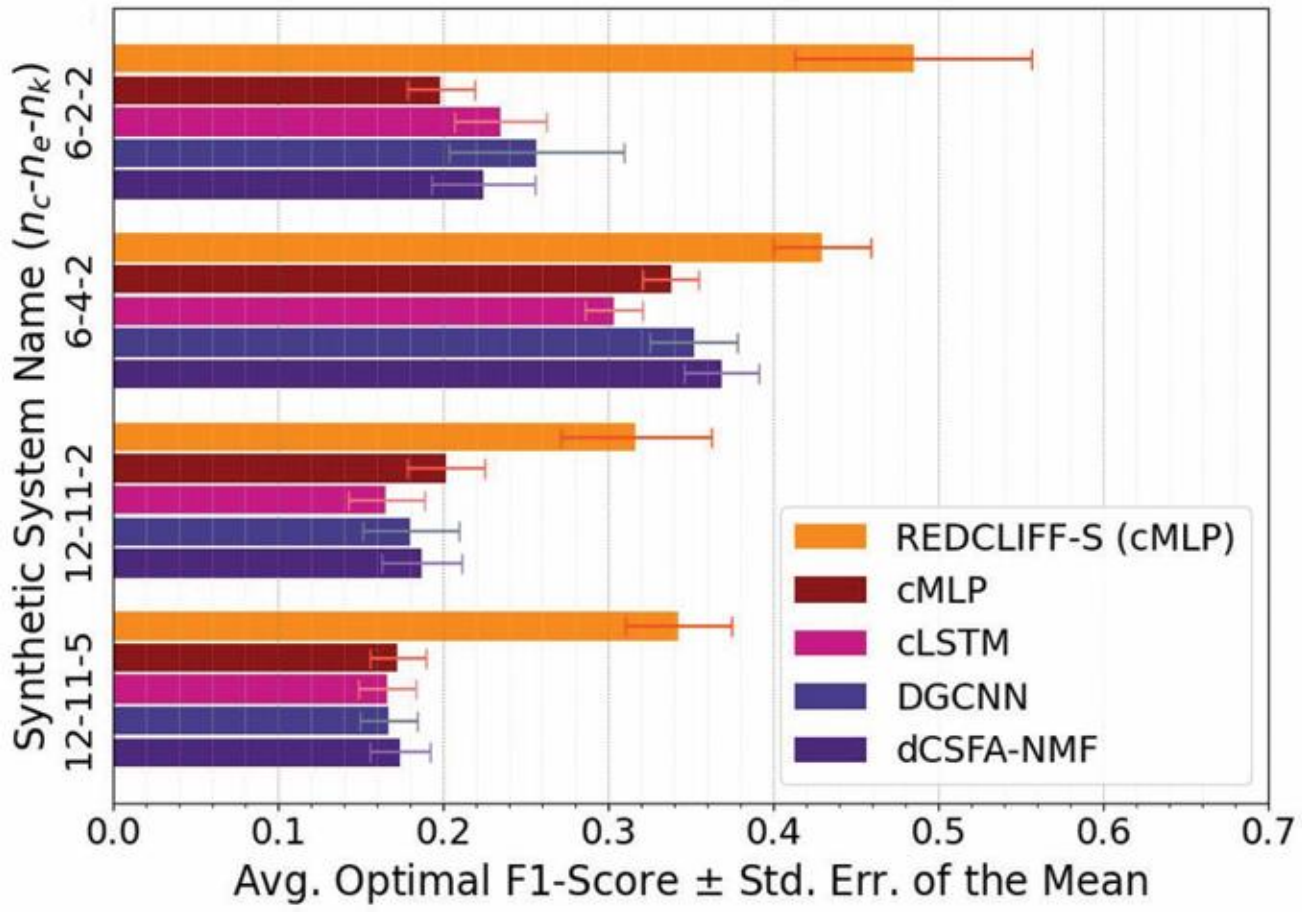

Example: REDCLIFF-S predictions on the 12 node, 11 edge, 5 factor system. Note the D-G-F path is similar to the true D-C-A-F path.

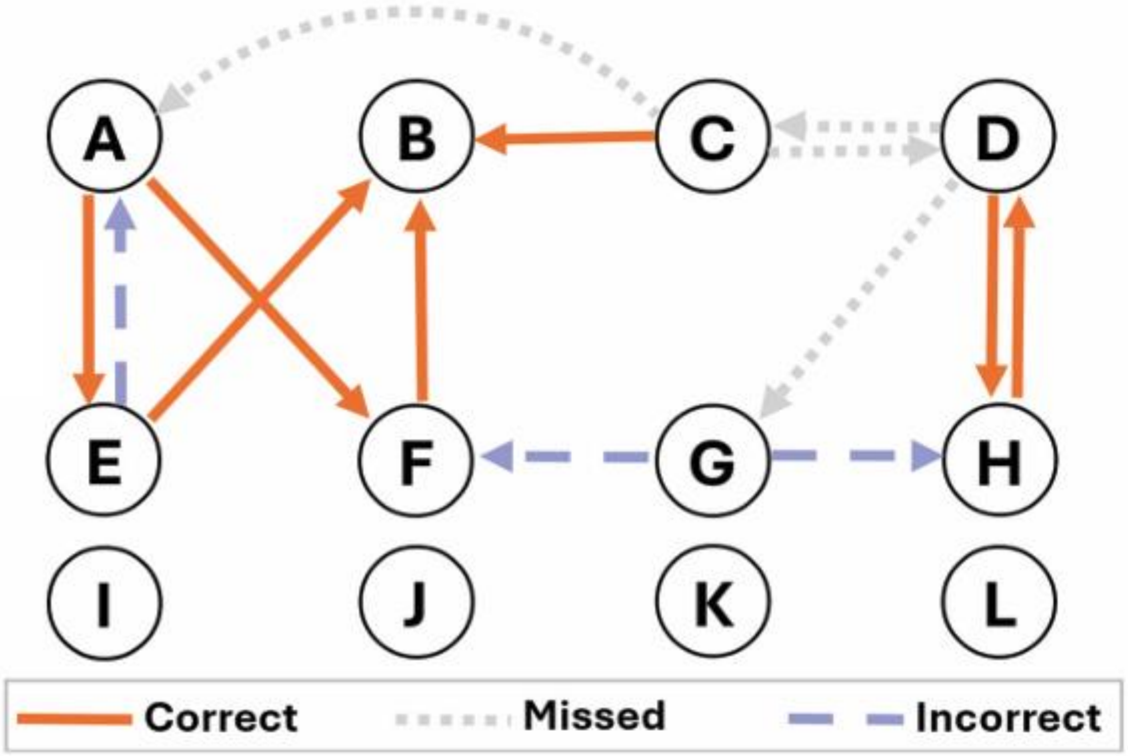

Experiments on Synthetic Data: Highlights

Below, we chart the mean performance of REDCLIFF-S and several supervised algorithms in generating hypotheses for the 12 node, 11 edge, 2 factor Synthetic System - across 10 random initializations, factors, and repeats - based on several accuracy and causal similarity metrics. **Performance improves with circle size (the bigger, the better).**

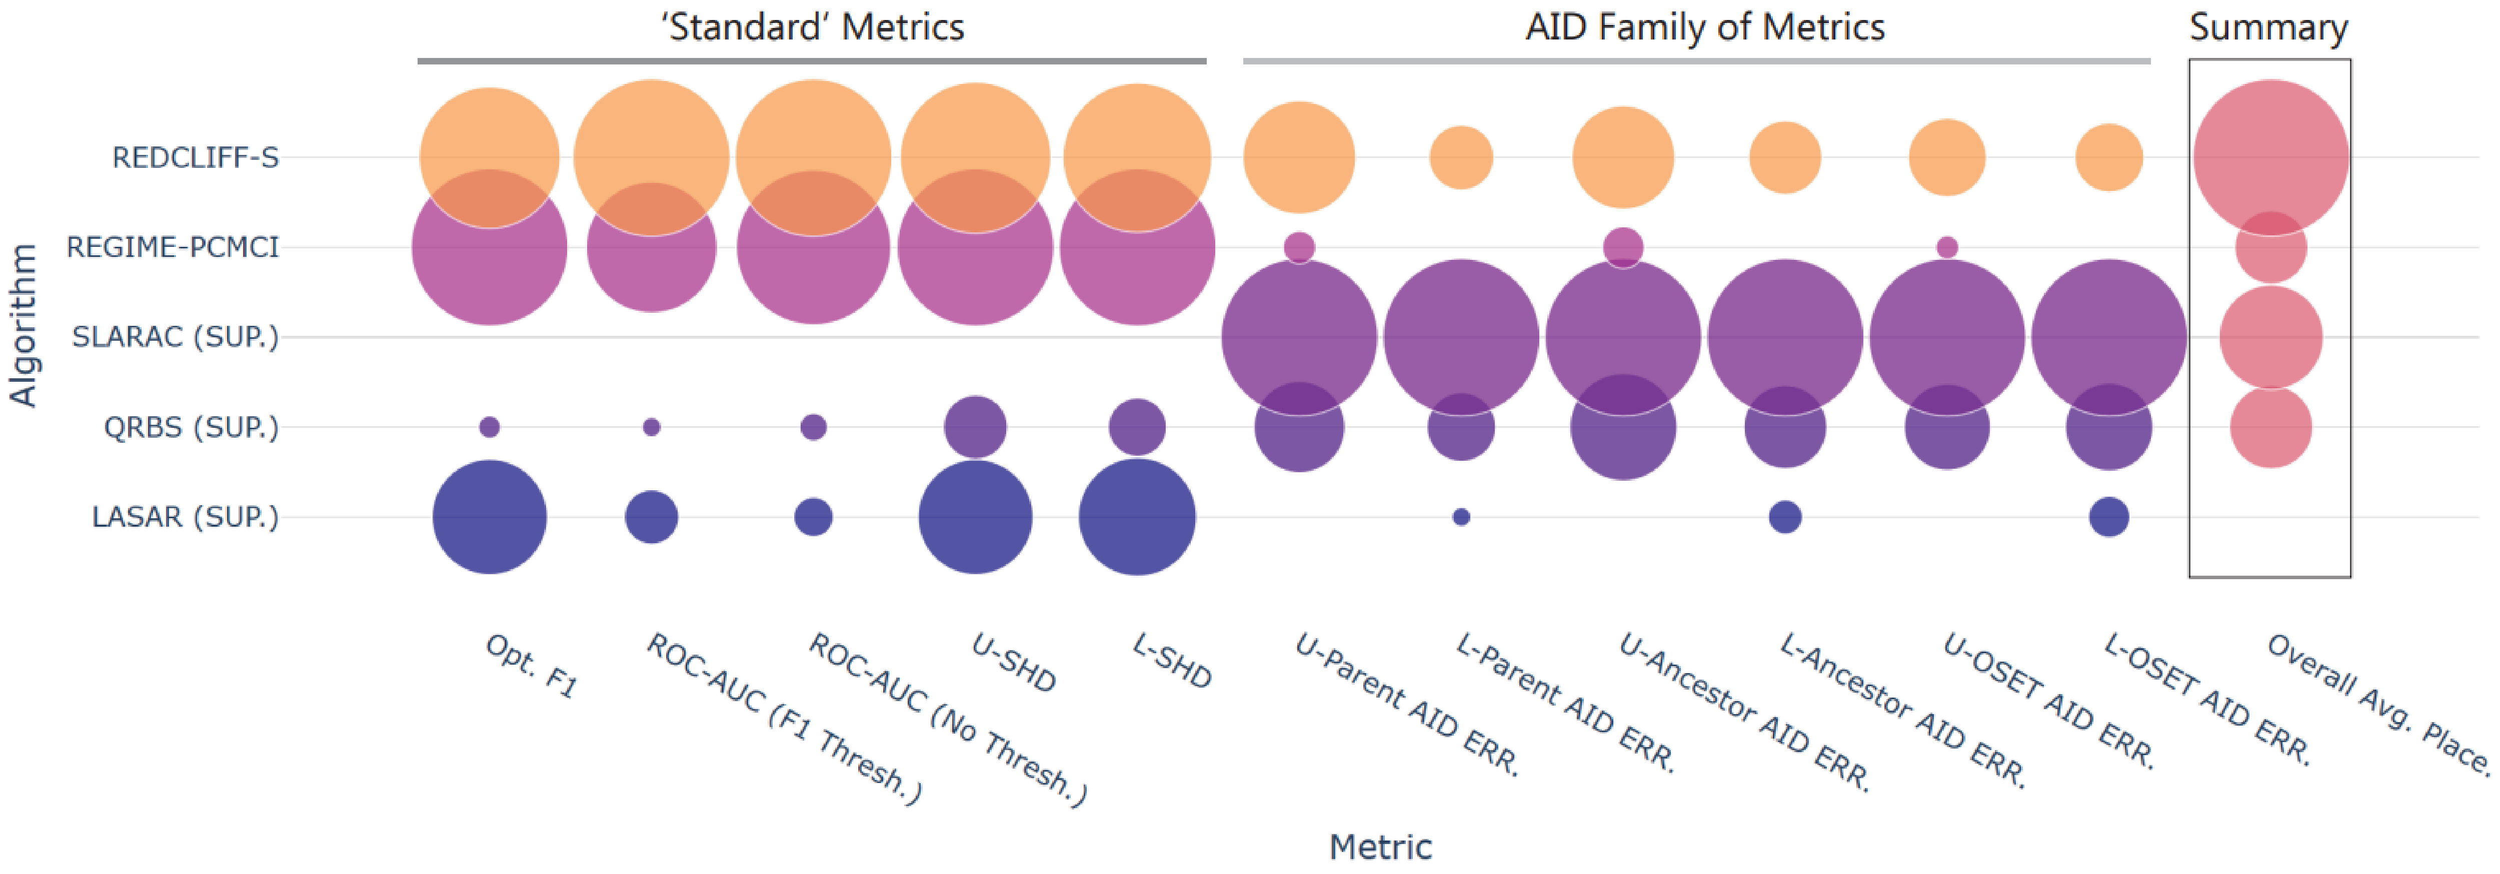

Findings on Real-World Mouse Model Data

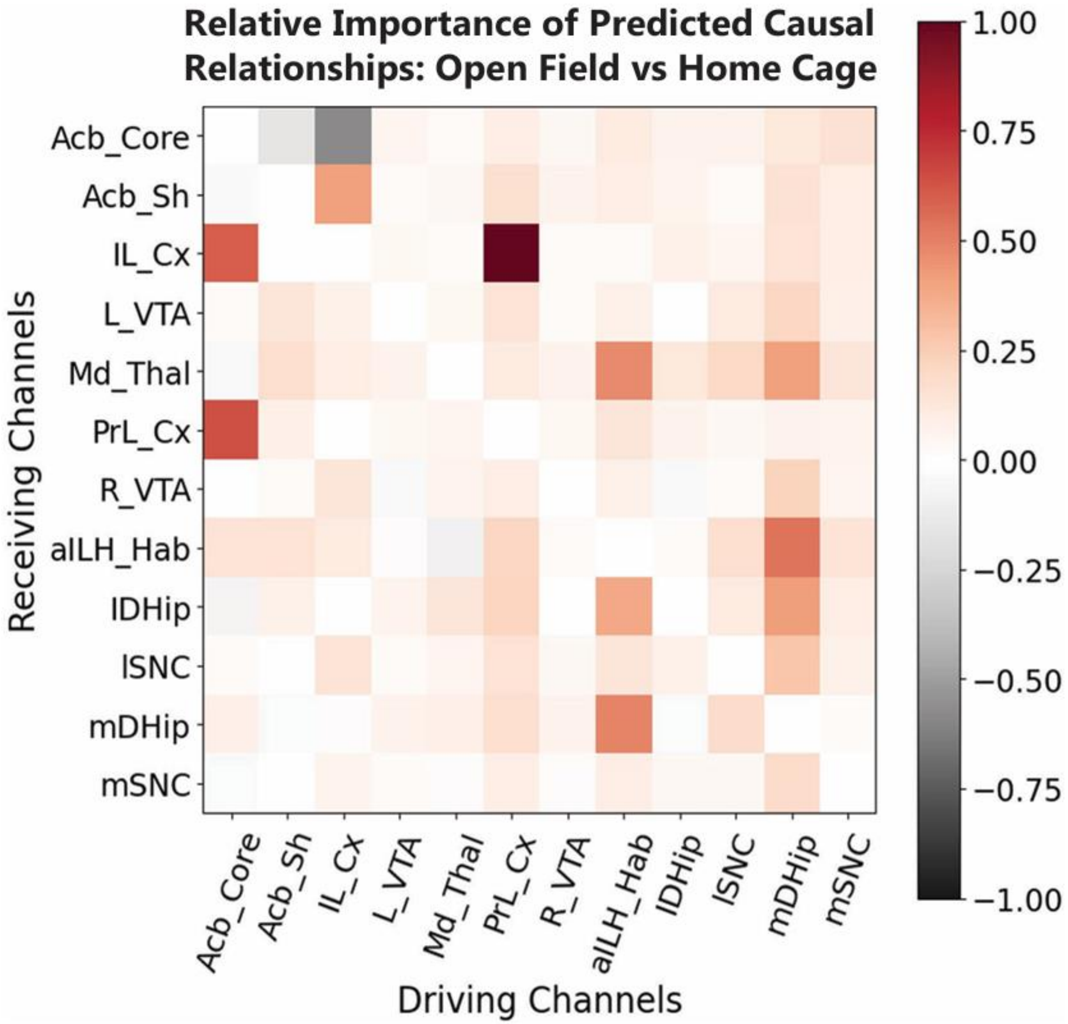

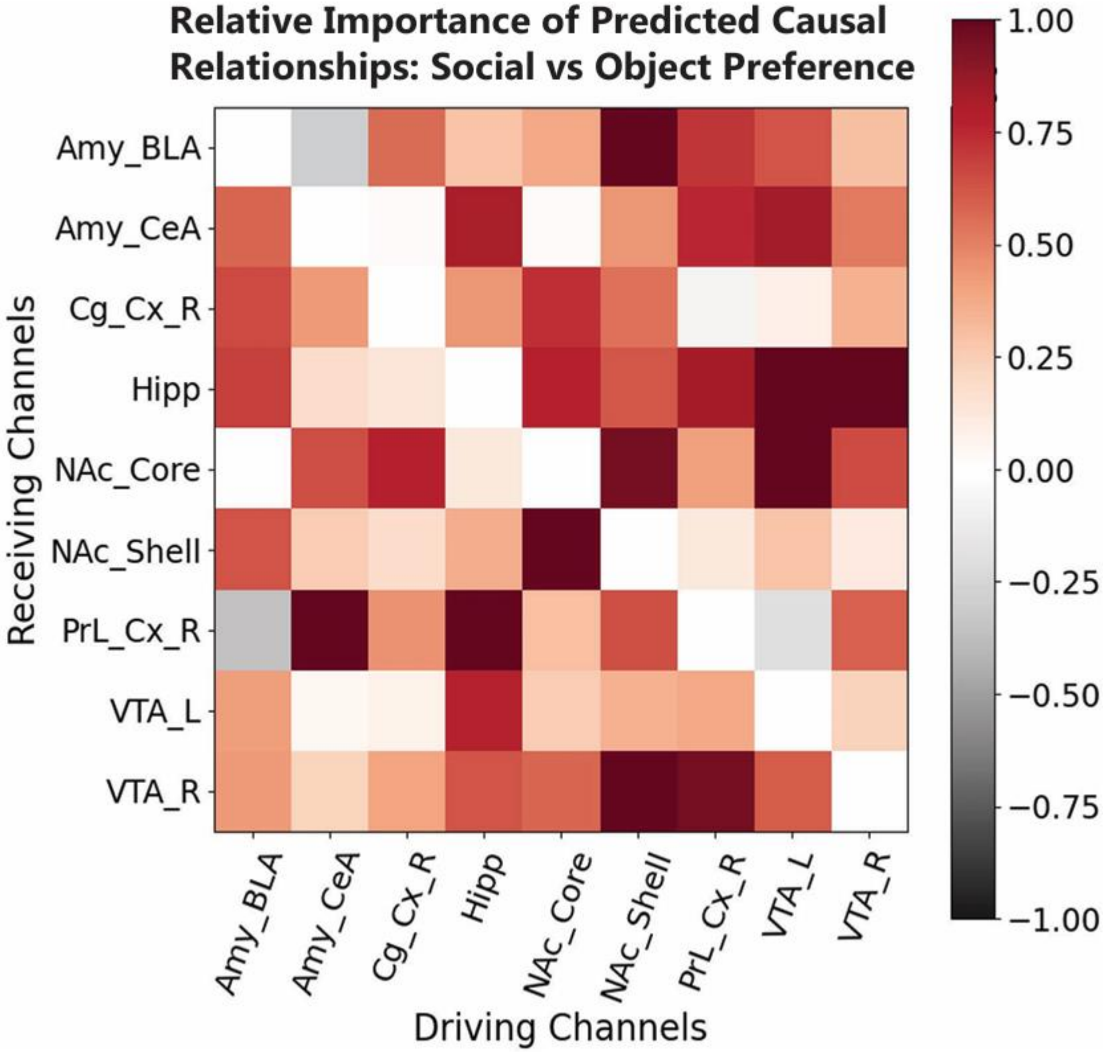

Funding and Additional Resources

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 2139754. Research reported in this publication was also supported by the National Institute Of Mental Health of the National Institutes of Health under Award Number R01MH125430. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Paper (Arxiv)

Code Repository

LinkedIn