

Estimation and validation of individualized dynamic brain models with resting state fMRI



Matthew F. Singh^{a,b,c,*}, Todd S. Braver^b, Michael W. Cole^d, ShiNung Ching^{c,aff5^e}

^a Department of Neuroscience, Washington University in St. Louis, St. Louis, MO, USA

^b Department of Psychology, Washington University in St. Louis, St. Louis, MO, USA

^c Department of Electrical and Systems Engineering, Washington University in St. Louis, St. Louis, MO, USA

^d Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, NJ, 07102, USA

^e Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, MO, USA

ARTICLE INFO

Keywords:

Resting state fMRI
Neural dynamics
Causal modeling
Recurrent neural networks
Dynamic functional connectivity

ABSTRACT

A key challenge for neuroscience is to develop generative, causal models of the human nervous system in an individualized, data-driven manner. Previous initiatives have either constructed biologically-plausible models that are not constrained by individual-level human brain activity or used data-driven statistical characterizations of individuals that are not mechanistic. We aim to bridge this gap through the development of a new modeling approach termed Mesoscale Individualized Neurodynamic (MINDy) modeling, wherein we fit nonlinear dynamical systems models directly to human brain imaging data. The MINDy framework is able to produce these data-driven network models for hundreds to thousands of interacting brain regions in just 1–3 min per subject. We demonstrate that the models are valid, reliable, and robust. We show that MINDy models are predictive of individualized patterns of resting-state brain dynamical activity. Furthermore, MINDy is better able to uncover the mechanisms underlying individual differences in resting state activity than functional connectivity methods.

1. Introduction

To understand human brain function, it is necessary to understand the spatial and temporal computations that govern how its components interact. This understanding can take multiple levels, ranging from statistical descriptions of correlations between brain regions to generative models, which provide a formal mathematical description of how brain activity evolves in time. However, efforts have taken quite different approaches based upon what data is available in human vs. nonhuman subjects. Several international neuroscience initiatives have relied upon nonhuman subjects to collect vast amounts of anatomical and electrophysiological data at the cellular scale (Markram, 2006, Markram et al., 2011, Okano et al., 2016). Generative models are then formed by integrating these cellular-level observations with known neuronal biophysics at the spatial scale of individual neurons or small populations (Markram, 2006, Markram et al., 2011).

In contrast, another set of large initiatives has instead focused on modeling individual human brain function using an approach often referred to as “connectomics” (e.g., Human Connectome Project, (Essen et al., 2013)). This approach relies on descriptive statistics, typically

correlation between fluctuating activity signals in brain regions assessed during the resting state (“resting state functional connectivity” or rsFC; (Biswal et al., 1995)). As a result, it is sometimes difficult to make mechanistic inferences based upon functional connectivity correlations (Buckner et al., 2013). Moreover, neural processes are notoriously nonlinear and inherently dynamic, meaning that stationary descriptions, such as correlation/functional connectivity, may be unable to fully capture brain mechanisms. Nevertheless, rsFC remains the dominant framework for describing connectivity patterns in individual human brains.

Despite the promise of human connectomics, there have been only a few attempts to equip human fMRI studies with the sorts of generative neural population models that have powered insights into non-human nervous systems. Notable advances have occurred in direct-parameterization approaches, with methods being developed to identify directed, causal influences between brain regions (e.g. Razi et al., 2017). Conversely, neural mass modeling approaches have also been extended to study human brain activity in a generative fashion (Breakspear, 2017), and these have provided new insights into the computational mechanisms underlying fMRI and MEG/EEG activity dynamics

* Corresponding author. 600 S. Euclid Ave., 63110, St. Louis, MO, USA.

E-mail address: f.singh@wustl.edu (M.F. Singh).

<https://doi.org/10.1016/j.neuroimage.2020.117046>

Received 20 December 2019; Received in revised form 8 May 2020; Accepted 8 June 2020

Available online 27 June 2020

1053-8119/© 2020 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

(Sanz Leon et al., 2013, Schirner et al., 2015, Jirsa et al., 2017, Demirtaş et al., 2019, Wang et al., 2019). However, unlike (linear) data-driven approaches (e.g. Dynamic Causal Modeling; Friston et al., 2003, Razi et al., 2017), neural mass models have been limited to replicating higher-level statistical summaries, such as functional connectivity, rather than predicting the actual time-series. This fact may not be relevant for some applications in which statistical descriptions will suffice. However, there remain many applications in basic neuroscience, neural medicine, and neural engineering for which more precise descriptions could be profitably leveraged.

Unfortunately, current approaches of both types have important limitations. In particular, the existing approaches to directly parameterize models (e.g. DCM) are subject to potential misinferences due to assumptions of linearity (Tu et al., 2018), and, in some cases, limitation to a relatively small number of brain regions (Ryali et al., 2011, Roebroek et al., 2011, Lohmann et al., 2012). This number has increased dramatically in recent years by assuming a fixed hemodynamic response function (Frässle et al., 2017), but remains well below modern brain parcellations, which feature several hundred regions (e.g. (Glasser et al., 2016, Schaefer et al., 2017), See Discussion). Likewise, with current neural mass modeling approaches, their ability to quantitatively recreate key features of individual-level functional connectivity has also been limited (Honey et al., 2009, Demirtaş et al., 2019, Wang et al., 2019). This may be because the most common approach is to parameterize connectivity from estimates of white matter integrity from diffusion imaging, which also can lead to potential misinference, since these connectivity estimates are constrained to be symmetric and positive ((Knock et al., 2009)). Efforts have been made to personalize these models by using individualized diffusion imaging data rather than group-average and/or tuning a small number of free-parameters to better approximate each subject's summary statistics (e.g. (Schirner et al., 2015, Jirsa et al., 2017, Demirtaş et al., 2019)). However, again, these models are not directly inferred from the brain activity time-series, which could limit their ability to accurately simulate the dynamical features of these time-series. Indeed, up until this point, it has not been shown that individual-level brain models can be directly parameterized and fit from fMRI while retaining sufficient complexity to capture – and predict – whole-brain activity. This limitation is critical because in order to accurately characterize individual variation in humans – which is the goal of personalized neuroscience and precision medicine initiatives (Ashley, 2015, Psaty et al., 2018, Satterthwaite et al., 2018) – individualized whole-brain models are required.

In the current work, we aim to fill this gap, by advancing high-resolution characterization of the human connectome through the parameterization of nonlinear dynamical systems models that go beyond statistical correlation matrices. The models consist of hundreds of interacting neural populations, each of which is modeled as an abstracted neural mass model evolving over time-scales commensurate with fMRI. Most critically, the models are optimized to capture brain activity dynamics at the level of individual human subjects. We present a computationally efficient algorithm to rapidly fit these models directly from human resting-state fMRI. The algorithm extends data-driven techniques towards the estimation of biologically interpretable models, and conversely enables the parameterization of dynamical neural models in a data-driven, individualized fashion with relatively few priors on the dynamics within and between brain regions. Our approach represents a significant departure and alternative approach to that of previous modeling efforts, in that every parameter in our model is individually estimated without consideration of prior anatomical constraints or long-term summary statistics.

We describe our efforts to develop and validate these models, demonstrating that they successfully characterize whole-brain activity dynamics at the individual level, and as such can be used as a powerful alternative to rsFC, and even to more closely related modeling approaches, such as DCM. Because of this goal, we term our modeling approach MINDy: Mesoscale Individualized Neural Dynamics. In the sections below, we introduce the MINDy modeling framework, highlighting its most

innovative and powerful features, and presenting results that validate its utility as an analytic tool for investigating the neural mechanisms and individual differences present in fMRI data.

2. Methods

2.1. Nature of interpretations from the model

The key premise of our approach is an expansion of the architectural description of brain networks from a simple connectivity matrix, to an interpretable dynamical model:

$$\dot{x} = W\psi_{\alpha}(x_t) - Dx_t + \varepsilon_t. \quad (1)$$

This model, which resembles a neural mass model (Wilson and Cowan, 1972, Hopfield, 1984, Deco et al., 2011, Breakspear, 2017) describes the evolution of brain activity at each anatomical location (each element of the vector x_t). Unlike true neural-mass models, we model abstracted brain activity commensurate with the fMRI timescale, rather than the evolution of population firing rate over milliseconds. Our model is similar, however, in that it is described by three components: a weight matrix (W) which identifies pathways of causal influence between neural populations, a parameterized sigmoidal transfer function (ψ) which describes the relation between the local activity of a population and its output to other brain regions (Eq. (3), Fig. 1 A, (Marreiros et al., 2008)), and a diagonal decay matrix (D) which describes how quickly a given neural population will return to its baseline state after being excited (i.e. the time-constant; Fig. 1 A). Process noise is denoted ε_t and is assumed to be uncorrelated between parcels. The additional parameters (α and D) reflect regional variation in intrinsic dynamics (D) and efferent signaling (α); critically, as described below, these parameters also show consistent anatomical distributions. These properties vary with brain network and are consistent even at the finer within-network scale (Fig. 4 A and B). Thus, our model, like a neural mass model, parameterizes both the interactions between brain regions and the processes that are local to each brain region that make it distinct.

It is important to recognize that this model is a phenomenological model in the sense that the state variables are more abstract than encountered in traditional mean-field models which combine biophysical first-principles and phenomenological approximations (e.g. the sigmoidal nonlinearity). Thus, inferences gained from the model are bounded by the inherent limitations of fMRI data (e.g. low temporal resolution and the indirectness of BOLD). The parametric form that we have chosen leads itself to interpretability. However, we stress that interpretability should not be confused with biophysical equivalence. As described in SI, there are likely many biophysical processes (including non-neuronal) contributing to each estimated parameter (5.1).

2.2. Robust estimation of individualized neural model parameters

While theoretical neural mass models operate in continuous-time, fMRI experiments have limited temporal sampling rates. Therefore, we approximate the continuous time neural model by fitting a discrete-time analogue for temporal resolution Δt (e.g. the sampling TR; Fig. 1 B):

$$x_{t+\Delta t} - x_t = (W\psi_{\alpha}(x_t) - Dx_t + \varepsilon_t)\Delta t. \quad (2)$$

Parameter estimation in the MINDy algorithm contains three main ingredients, which ensure that estimates are robust, reliable, and valid. First, the transfer functions of neural mass models are allowed to vary by brain region through the scalar parameter α :

$$\psi_{\alpha}(x_t) := \sqrt{\alpha^2 + (bx_t + .5)^2} - \sqrt{\alpha^2 + (bx_t - .5)^2}. \quad (3)$$

Each brain region has its own α parameter, fit on a subject-wise basis, while b is a fixed global hyperparameter ($b = 20/3$ for the current case). The use of a parameterized sigmoid allows for additional anatomical

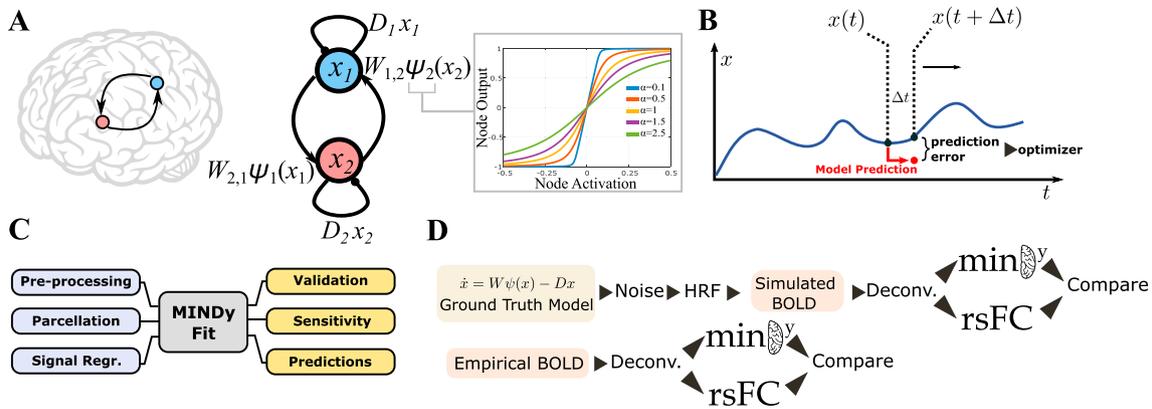


Fig. 1. Overview of Methods Employed. A) The MINDy model consists of coupled 1-dimensional neural-mass models (Hopfield form (Hopfield, 1984)). The shape of the transfer-function for each brain region is parameterized by a curvature parameter α . B) Model goodness-of-fit was measured through one-step prediction of the empirical time-series. C) Overview of data processing and analyses: data was processed according to Siegel and colleagues (Siegel et al., 2017) and parcellated. Reported analyses fall into three categories: validation, sensitivity to nuisance parameters, and predictions of brain activity patterns. D) In both simulations and empirical analyses the BOLD signal was Wiener-deconvolved (Weiner, 1949) with a canonical HRF function (see Methods; (Friston et al., 1998)) before being analyzed with either MINDy or rsFC.

heterogeneity in region-wise dynamics. This form of transfer function is general enough to capture conventional choices (see SI Sec. 5.2 for a derivation of the function and its relation to conventional transfer functions). Secondly, we make use of recent advances in optimization to ensure that the fitting procedure (SI Fig. 9 C) is robust. By using Nesterov-Accelerated Adaptive Moment Estimation (NADAM, (Dozat, 2016)) we achieve the speed advantage of stochastic gradient descent (SGD) algorithms, while at the same time preventing both over-fitting and under-fitting (see SI Sec. 5.3 for discussion). This approach leads to a very reasonable time duration for estimation (approximately 1 min on a standard laptop; see SI Sec. 5.8 for a comparison with spDCM).

Lastly, we constrain the problem by decomposing the large matrix of connection weights (W) into two simultaneously fit components: a sparse component W_S and a low-dimensional component $W_L := W_1 W_2^T$ in which both W_1 and W_2 are $n \times k$ rectangular matrices with n being the number of neural masses (brain parcels) and $k < n$ being a global constant that determines the maximum rank of W_L . This decomposition is advantageous for concisely representing the interactions of structured networks and is the most important element of the fitting process. Sparseness criteria were achieved through L_1 regularization (Donoho, 2006) with the resultant fitting objective being to minimize:

$$J = \frac{1}{2} E_T [\| (X_{T+\Delta t} - X_T) - [(W_S + W_L)\psi_\alpha(X_T) - DX_T] \|_2^2] + \lambda_1 \|W_S\|_1 + \lambda_2 \text{Tr}(|W_S|) + \lambda_3 (\|W_1\|_1 + \|W_2\|_1) + \frac{\lambda_4}{2} \|W_L\|_2^2 \quad (4)$$

The notation E_T denotes the mean over all temporal samples considered (the “minibatch” of each iteration) so the first term simply corresponds to the mean square error of predictions. Each of the remaining penalty terms have a global regularization constant (λ_i) that is shared across all subjects. This regularization scheme was adopted in order to reduce the dimensionality of the parameter estimation problem, while at the same time, attempting to reflect the consistently observed community-structure of brain connectivity measures. Under this view, brain connectivity patterns can be described in terms of communities (sub-networks) linked together by highly connected hubs. We envision the sparse component of connectivity to mimic the communication between connectivity hubs. By contrast the low-rank component is meant to account for the propagation of signals from hubs to their corresponding subnetworks and vice-versa.

We employ this two-component weight formulation as a heuristic that facilitates high-dimensional model fitting. In most analyses we only analyze the composite weight matrix rather than its components. However, preliminary results indicate that properties of this decomposition,

namely the ratio of sparse vs. low-rank components, may be a marker of individual differences (see SI Sec. 5.6). Interestingly, recent work by Mastrogioiuseppe and Ostojic (2018)) has also considered models in which connectivity is the sum of two terms: one low-rank and one random. The authors found that these structures produced low-dimensional dynamics which could be predicted based upon network structure and exogenous (task) input. Such analyses may be relevant for understanding the role of connectivity in MINDy. Bayesian and algebraic interpretations of this penalty function are presented in SI Sec. 5.5. We also discuss the well-posedness of this problem (SI Sec. 5.5).

Throughout, we use the term “weights” to refer to the matrix W in estimated dynamic neural models. This is to differentiate the model connectivity parameter from the term “resting-state functional connectivity” (rsFC), which instead refers to the correlation matrix of BOLD time-series, rather than the mechanistic concept that it is often assumed to measure (i.e. direct and indirect interactions between brain regions). We reserve the term “effective connectivity” to indicate a causal, monotone relationship in activity between brain regions that evolves over no more than 2s (the typical fMRI sampling rate). Thus, both the fit model weights and the rsFC are ways to approximate the effective connectivity, even though rsFC may not support reverse inferences regarding directedness and causality.

2.3. Study design

The objective of the current study was to rigorously validate a new approach for data-driven whole-brain modeling (MINDy). The study design consisted of both numerical simulations to validate the accuracy of models with respect to a known ground-truth, as well as empirical analyses of HCP resting-state data. The latter analyses were designed to test whether MINDy adds additional value in-practice and to quantify its performance in the presence of known experimental confounds (e.g. motion).

2.4. Empirical dataset

2.4.1. HCP resting-state scans

Data consisted of resting state scans from 53 subjects in the Human Connectome Project (HCP) young adult cohort, 900 subject release (for acquisition and minimal preprocessing details, see (Glasser et al., 2013); WU-Minn Consortium). Each subject underwent two scanning sessions on separate days. Each scan session included two 15-min resting-state runs (two scans \times two days) for a total resting state scan time of 60 min (4800 TRs). The two runs for each session corresponded to

acquisitions that had left-right and right-left phase-encoding directions (i.e., balanced to account for potential asymmetries in signal loss and distortion). The TR was 720 ms and scanning was performed at 3T. The subjects were selected by starting with an initial pool of the first 150 subjects and then excluding subjects who had at least one run in which more than 1/3 of frames were censored (i.e. 400 bad frames out of 1200).

Although this criterion greatly decreased the number of useable subjects from the initial pool of 150 to 53 (attrition = 65%), it should be noted that it is likely to be overly conservative. We employed such a strongly conservative criterion for this first-stage validation effort to provide the cleanest data from which to test the model. Likewise, we had the luxury of drawing upon a very large-sample dataset. In contrast, we believe that the exclusion criteria will not need to be as conservative in a research setting for which model cross-validation is not performed on every subject (i.e., it is probably overly stringent to require that all four sessions be clean, since we only used two sessions at a time). In particular, the use of cross-validation required that two models be fit for every subject using disjoint data so that the validation required twice as much data as would normally be required. Moreover, we required that the data be uniformly clean so that we could parametrically vary the amount of data used (i.e. criteria were in terms of absolute cleanness for each scanning session rather than number of clean frames). However, there is no reason why the models could not be fit to clean segments of scanning sessions.

2.4.2. Preprocessing

Data were preprocessed through the rsFC pipeline proposed by Siegel and colleagues (Siegel et al., 2017; SI Fig. 9 A). The first stage of this pipeline is the HCP minimal pre-processing pipeline (see (Glasser et al., 2013)) with FSL's ICA-FIX correction (Griffanti et al., 2014, Salimi-Khorshidi et al., 2014). We then applied one of 3 s-stage pipelines developed by Siegel and colleagues (Siegel et al., 2017); Sec. 3.7.3), to test the effects of including various additional preprocessing steps. In all three pipelines, drift was mitigated by detrending data. The pipelines also all included motion scrubbing, using both Framewise Displacement (FD) and the temporal derivative of variation (DVARs). Frames that exceeded the cutoffs for FD (.2 mm) or DVARs (5% above median) were replaced via linear interpolation (Power et al., 2015). Respiratory artifact was mitigated with a 40th-order 0.06-0.14 Hz band-stop filter applied to FD and DVARs for all pipelines (Siegel et al., 2017).

The 3 s-stage pipeline variants differed however, in the number of regressors included to remove nuisance signals. The first variant mainly corrected frame-to-frame motion artifact, which has been found to induce systematic errors in functional connectivity studies, i.e. generating spurious short-distance correlations while diminishing long distance ones (Power et al., 2012). In addition to data scrubbing, motion correction was performed using the 12 HCP motion regressors and their temporal derivatives. The second, more extensive pipeline variant, known as CompCor, also removed cardiac and respiratory signals, by additionally regressing out principal components of the white matter and cerebrospinal fluid signals (Behzadi et al., 2007). Lastly, the third pipeline variant also added global signal regression (GSR; (Aguirre et al., 1998a)), in which the mean signals from white matter, cerebrospinal fluid, and grey matter are also included as regressors. As the variables included are cumulative, these three pipelines form a representative hierarchy of preprocessing approaches, that optionally includes CompCor or CompCor + GSR in addition to motion scrubbing. For most analyses we used the full (third) pipeline, but we also compared the effects of pipeline choice (Sec. 3.7.3).

After the second-stage preprocessing pipelines, we deconvolved the parcellated data (see below) with the generic SPM hemodynamic kernel (Friston et al., 1998) using the Wiener deconvolution (Weiner, 1949). For the Weiner deconvolution, we used noise-power to signal-power parameter .02. The value of this parameter dictates the degree of temporal filtering during the deconvolution with smaller values being more

parsimonious (less additional filtering). We then smoothed by convolving with the [0.5 0.5] kernel (2 point moving average) and z-scored the result. To test the robustness of the fitting procedure, we compared the effect of the second-stage preprocessing pipelines for some analyses. Based upon these results, we chose the third variant pipeline (GSR + CompCor + motion) for all other analyses. For all empirical rsFC analyses we use the deconvolved data to prevent bias from the deconvolution procedure in comparing MINDy and rsFC. As described further below, we also tested the effect of mismatches between "true" and canonical HRF models (Sec. 3.7.4, 3.7.5).

We defined derivatives in terms of finite differences. Since HCP employed unusually fast scanner TRs, we temporally downsampled the estimated derivatives for calculating goodness-of-fit in non-simulation analyses to represent the anticipated benefits to typical fMRI protocols and improve SNR: $dX(t) = (X(t + 2) - X(t))/2$.

2.4.3. Parcellation atlases

In the present framework we define whole-brain models in terms of connected neural populations. Thus, the approach demands that the neural populations be defined a-priori. For the present case of fMRI data, we define these populations to be anatomical brain regions corresponding to subcortical structures and cortical parcels. For subcortical regions, we follow the HCP protocol in considering 19 subcortical regions as defined by FreeSurfer (Fischl, 2012). For cortical parcels, we generally employed the gradient-weighted Markov Random Field (gwMRF) parcellation with 200 parcels per hemisphere (Schaefer et al., 2017) and organized according to the 17 cortical networks described in (Thomas Yeo et al., 2011). The gwMRF parcellation is optimized to align with both resting-state and task fMRI, and has been found to demonstrate improved homogeneity within parcels relative to alternative parcellation techniques. However, for anatomical analyses we compared with an additional atlases (SI Fig. 11 C,G) to ensure generality: the MMP atlas (Glasser et al., 2016) which was also derived from a combination of rest and task-based data. The MMP (Multi-Modal Parcellation) atlas is symmetric with 180 parcels per hemisphere.

2.5. MINDy fitting procedure

MINDy models were fit by applying the iterative NADAM algorithm (Dozat, 2016) to optimize the MINDy cost-function (Eq. (4); see SI Sec. 5.12). This algorithm belongs to the family of stochastic gradient-descent techniques and we provide further detail/discussion regarding NADAM in SI Sec. 5.3. To ensure algorithmic stability, we used two transformations (one each for the curvature and decay parameters) which are detailed in SI Sec. 5.12. The gradient equations for each parameter in detailed in SI Tab. 11.

2.5.1. Compensating for regularization bias

In order to retrieve parsimonious weight matrices and reduce overfitting, we employed regularization to each weight matrix (both the sparse and the low-rank matrices) during the fitting process. One consequence of regularization, however, is that the fitted weights may be unnecessarily small as weight magnitudes are penalized. After fitting, we therefore performed a global rescaling of weight and decay contributions for each model using robust regression (Holland and Welsch, 1977) as implemented by MATLAB2018a. Specifically, we fit two scalar parameters: p_W, p_D in regressing $dX(t) = p_W W \psi(X(t)) - p_D D x$ collapsed across all parcels. Here p_W and p_D represent global rescaling coefficients for the weights and decay, respectively. As this compensating step only used global rescaling for W and D , it had no effect upon the relative values for each parcel, only the total magnitude of the W and D components. Since only two values are estimated, this step does not reintroduce overfitting. Although we performed this step using robust regression, we obtained identical results using conventional linear regression. The choice of robust regression was made as a safeguard for high leverage points as might occur due to motion artifact. However, results indicate that

conventional regression may suffice for sufficiently clean data.

2.5.2. Selecting hyperparameters and initialization

The proposed fitting procedure requires two sets of hyperparameters: the four regularization terms specific to our procedure and the four NADAM parameters (Dozat, 2016). By “hyperparameters” we refer to free constants within an algorithm which distinguishes them from the “parameters” of an individualized model.

Hyperparameters were hand-selected for model goodness-of-fit and reliability, based upon prior numerical exploration with a subset of 10 subjects who did not belong to the “data source” subjects. Thus, these subjects were not included in any further analyses so the hyperparameter selection procedure did not artificially inflate model performance. The selection criteria were to maximize cross-validated goodness-of-fit under the constraint that test-retest correlations were greater than 0.7 for all parameters. Regularization values were sampled with resolution 0.005. The chosen set of hyperparameters was then constant for all test subjects. Hyperparameter values and discussion are included in SI (Tables 12, 13). The initialization distributions for the algorithm were similarly selected using the same subjects and are included in the SI (Table 12). We explored the effect of hyperparameter choices on the sparsity of MINDY relative to rsFC and found that for any choice of regularization hyperparameters (even 0), the group-average MINDY weights are sparser than rsFC (SI Sec. 5.7).

2.6. Ground-truth simulations

2.6.1. Realistic whole-brain simulations

For the analyses of sensitivity and individual differences we generated new, synthetic individuals by randomly sampling neural mass model parameters from the parameter distributions estimated from the full dataset (i.e. $N = 53$ participants). The decay and curvature parameters (α , D) were independently sampled for each parcel from that parcel’s population distribution. The weight matrices, however, were sampled as a whole rather than sampling each individual connection as we found that the latter approach led to pathological behavior in simulations. For the robustness analyses, ground truth models were drawn from those fit to experimental sessions. The ground-truth models were simulated as stochastic differential equations ($dX = f(X)dt + \sigma_W dW_t$) with $f(X)$ the deterministic neural mass model and units time measured in terms of the fMRI TR. Models were Euler-Maruyama integrated with $dt = 1/4$ and $\sigma_W = .45$ in units TR (720ms) to generate simulated neural activity time-series. Neural-activity was then downsampled to 1 TR resolution (as opposed to the simulation’s time-step of $dt \times TR$) and convolved with the SPM-style HRF kernel (Friston et al., 1998); SI Fig. 20 C):

$$h(t; \alpha_{(1,2)}, \beta_{(1,2)}, c) := \frac{t^{\alpha_1-1} e^{-\beta_1 t} \beta_1^{\alpha_1}}{\Gamma(\alpha_1)} - \frac{t^{\alpha_2} e^{-\beta_2 t} \beta_2^{\alpha_2}}{c\Gamma(\alpha_2)} \quad (5)$$

Here Γ is the gamma function (equal to factorial for integer values). The parameters describe two gamma-distributions (one α, β pair per distribution) and a mixing coefficient (c) to generate a double-gamma distribution. Parameters were set to their default values ($\alpha_1 = 6, \alpha_2 = 16, \beta_1 = 1, \beta_2 = 1, c = 1/6$) except for the simulation featuring HRF variability. In this case, random perturbations were added to each parameter and were drawn from the normal distribution with mean zero and SD as indicated. The final simulated BOLD signal was then generated by adding white, gaussian noise with the indicated SD (Fig. 1 D).

2.6.2. Randomized network simulations

Although some ground-truth simulations leveraged the empirical MINDY distributions to maximize realism (Sec. 3.2.1, 3.2.2, 3.7.1, 3.7.4), others used randomly generated networks of Hopfield or neural mass models (Sec. 3.7.5, 3.8, SI Sec. 5.8). The latter ground-truth simulations prevent circularity (i.e. using MINDY distributions to test MINDY) by drawing parameters from random hyperdistributions independent of

previous analyses. These distributions were designed to possess complex network structures by superimposing three simpler network structures: community-structure (M_1), sparse structure (M_2), and low rank structure (M_3). These distributions are characterized by standard-deviation parameters σ_1 and σ_2 . An asymmetry parameter σ_a characterizes the degree to which the resultant network is asymmetric. Each standard-deviation parameters was randomly sampled for each ground-truth model from normal distributions: $\sigma_1, \sigma_a \sim N(4, .05^2)$ and $\sigma_2 \sim N(3, .05^2)$. Connectivity matrices were then randomly parameterized as follows:

$$\begin{aligned} M_1 &\sim \left[N(0, 1/\sigma_1^2) + N(0, 1/\sigma_1^2) \right]_{n/q \times n/q} \\ M_2 &\sim \left[N(0, 1/\sigma_2^2) \right]_{n \times n} \\ M_3 &\sim \left[N(0, 1/\sigma_1^2) + N(0, 1/\sigma_1^2) \right]_{n \times k} \times \left[N(0, 1/\sigma_1^2) + N(0, 1/\sigma_1^2) \right]_{k \times n} \end{aligned} \quad (6)$$

Here, the bracket outside each matrix denotes its size with $n = 40$ denoting the total number of nodes, q denoting the number of nodes per community (randomly set to either 1 or 2 with equal probability), and $k = 5$ denoting the rank of the low-rank component. We denote the Kronecker product \otimes and use it to copy the community level matrix (M_1) among each node belonging to the community: $\widehat{M}_1 := \mathbf{1}_{q \times q} \otimes M_1$. The three component matrices are then combined as follows:

$$Q = \widehat{M}_1 + M_2 + M_3, \quad \widehat{Q} = (Q + (Q - Q^T) / \sigma_a) \quad (7)$$

The final matrix C is formed by censoring elements of \widehat{Q} whose absolute value is below $1/4$ the standard deviation of \widehat{Q} . This same technique was used to randomly generate networks of Hopfield models with homogeneous, heterogeneous, or nonlinear hemodynamic effects and realistically-parameterized neural mass models with nonlinear hemodynamics.

2.6.3. Hopfield network simulations

We employed two cases of non-MINDY ground truths: Hopfield networks and neural-mass models (Sec. 3.7.5, 3.8). Continuous, asymmetric Hopfield models are similar in form to the MINDY model, but use a tanh transfer function:

$$dx = (W \tanh(b_0 \circ x) - Dx)dt + \sigma_W dW. \quad (8)$$

Here, the slope vector $b_0 \in \mathbb{R}^n \sim N(6, (.5)^2)$ and diagonal elements of the decay matrix D drawn from $N(4, (.1)^2)$ (non-diagonal elements are zero). As elsewhere, the symbol \circ denotes the Hadamard product (element-wise multiplication). Models were simulated via Euler-Maruyama integration with $dt = .1s, \sigma_W = 0.2, TR = .7s$, and total simulation length $t = 10,000$. We considered the case in which no hemodynamics are present, in which case MINDY is fed $x(t)$ downsampled according to TR, and the case in which $x(t)$ is convolved with spatially heterogeneous hemodynamics and deconvolved with the canonical HRF before being fit by MINDY. In the latter case, the HRF function was parameterized as before, but with the ground-truth α_1 parameter for each brain region drawn from $N(6, (.25)^2)$ and the β_1 parameter drawn from $N(1, (.25/6)^2)$. The simulated BOLD was produced by convolving the simulated time-series with the ground-truth HRF before temporal downsampling. In both cases, initial conditions for each node were independently drawn from $N(0, 1)$ and the first 100 samples were dropped. Since the total number of nodes was approximately one-tenth of those used in the HCP data, we rescaled the dimension of the low-rank component by one-tenth (from 150 to 15). Similarly, we rescaled the regularization terms inversely proportionate to the effect of rescaling W by a factor of 10: (λ_1, λ_3 by $1/10, \lambda_2$ by $1/\sqrt{10}$ and λ_4 by $1/10^2$). For simulations using the Balloon-Windkessel model of hemodynamics, $x(t)$ was rescaled to the range of average synaptic gating via $5S(t) = 1 +$

$\tanh(x(t)/10)$. This transformation of x was then substituted into the nonlinear hemodynamic model (below) to generate simulated BOLD signal. In all cases, time series were z-scored, smoothed via nearest-neighbor ([0.5 0.5] kernel) and run through MINDy for 150,000 iterations (approximately 70 s) with the original batch size of 250.

2.6.4. Neural mass and Windkessel-Balloon model simulations

For our neural mass ground-truth simulations (Sec. 3.7.5, 3.8), we largely followed the approach of Wang and colleagues (Wang et al., 2019) in using single-population neural mass models (20 masses/simulation in Sec. 3.7.5 and 6 to 16 in Sec. 5.8) with Windkessel-Balloon model hemodynamics (Friston et al., 2000). Similar to the MINDy model, the neural mass model (Hansen et al., 2015) contains a monotone nonlinearity ($\hat{\psi}$) and linear decay $1/\tau_S$:

$$\begin{aligned} \dot{S}_i &= -S_i/\tau_S + r(1 - S_i)H(x_i) + \sigma_w dW \\ H(x_i) &= \hat{\psi}(-d(ax_i - b)) \\ H(x_i) &= \hat{\psi}(-d(ax_i - b))x_i = \sum_j C_{ij}S_j + I_{Sub} \end{aligned} \quad (9)$$

The variable S describes the average synaptic gating, while H describes the population firing-rate. We used the default parameter settings: $\tau_S = .1s$, $a = 270n/C$, $b = 108Hz$, $d = .154s$, $r = .641$. Unlike Wang and colleagues (Wang et al., 2019), we used a logistic sigmoid transfer function for $\hat{\psi}(x) = 1/(1 + \exp(-x))$ instead of the rectified linear transfer function: $x/(1 + \exp(-x))$, as the former is less prone to pathological behavior in random networks. Subcortical input was $I_{Sub} = 5$. Connection weight matrices were randomly generated as described in the previous section, but with 1.5 added to all recurrent connections and the resultant matrix scaled by a factor of 100. Simulated neural activity is converted into BOLD signal through the Windkessel-Balloon model (Friston et al., 2000):

$$\begin{aligned} \dot{z}_i &= S_i - \kappa z_i - \gamma(f_i - 1) \\ \dot{f}_i &= z_i \\ \tau \dot{v}_i &= f_i - v_i^{\alpha_G^{-1}} \\ \tau \dot{q}_i &= \frac{f_i}{\rho} [1 - (1 - \rho)^{1/f_i}] - q_i v_i^{\alpha_G^{-1}-1} \end{aligned} \quad (10)$$

The variables z, f, v , and q model vasodilation, inflow, blood volume, and deoxyhemoglobin content, respectively. Parameters were: $\rho = .34$, $\kappa = .65s^{-1}$, $\gamma = .41s^{-1}$, $\tau = .98s$, $\alpha_G = .32$. The simulated BOLD signal at each TR is then modeled as:

$$BOLD(v_i, q_i) = V_0 \left[k_1(1 - q_i) + k_2 \left(1 - \frac{q_i}{v_i} \right) + k_3(1 - v_i) \right] \quad (11)$$

Resting blood volume fraction is denoted $V_0 = .02$. Scanning parameters k_1, k_2, k_3 were set to 3T values according to Demirtas and colleagues (Demirtas et al., 2019): $k_1 = 3.72, k_2 = .53, k_3 = .53$. Simulations were run with $dt = 25$ ms and $\sigma_w = .005$ for total length $t = 40,000$. Sampling was performed every 29 time-steps ($TR = 725ms$) and the first 10% of samples were dropped. The resulting time-series were deconvolved with the canonical HRF assumed by MINDy and z-scored. MINDy hyperparameters were identical to the rate-model case and MINDy was run for 10,000 iterations (approximately 6 s) with batch size 250. Initial conditions for hemodynamic variables were randomly sampled from $|N(0, 1)|$. Initial conditions for the neural variable (S) were generated by first sampling $S_0 \sim |N(0, 1)|$ and then performing the transformation $S_0/(1 + S_0)$.

2.7. Simulations for DFC analysis

For analyses of dynamic functional connectivity, models were estimated for each subject (one per session) using the full HCP temporal resolution $dX(t) = X(t + 1) - X(t)$. These models were then used to generate simulated resting-state fMRI data, but with additional process

noise added as would be expected in observed fMRI timeseries data. We used the same time-scale for simulation as in the validation models ($dt = .5$ TR). However, whereas the validation simulations employed process noise containing constant variance across parcels, we used a naive estimate of process noise for each parcel, that was based upon the residual error of model fits over subsequent time-steps. We avoided doing this in the validation stage so that ground-truth parameters could not be recovered simply by observing noise. The residual error covaried with the decay parameter across parcels at the group-level, but not at the individual level, despite individual differences in both noise and decay being reliable within parcel. We reintroduced parcel-based variation into the DFC simulations to obtain maximum realism. We considered both the case in which process noise was allowed to vary by parcel but not by individual within a test-retest group (e.g. using the mean noise across subjects for each session separately), as well as the case in which process noise was determined on a subject-wise basis. Results obtained with either method were near-identical for the DFC reliability analyses so we present results using the session-wise group-mean process noise (e.g. the mean process noise for each parcel averaged across all day 1 scans or all day 2 scans). Initial conditions were drawn from each subject's observed data for that scanning session. Simulations were run for 2600 time steps (1300 TRs) using 15 different initial conditions per session and temporally downsampled back to the scanning TR. After simulation, we downsampled from the 400 parcel to the 100 cortical parcel variants of gwMRF (Schaefer et al., 2017) and removed subcortical ROIs in order to reduce computational complexity of subsequent DFC analyses.

2.8. DFC analyses

DFC analyses consisted of the standard deviation and excursion (Zhang et al., 2018) of the time-varying correlation between brain regions. To calculate time-varying correlations we used Dynamic Conditional Correlation (DCC; (Engle, 2002)). To avoid confusion with other references to "standard-deviation" we refer to this measure as " σ -DFC" as it pertains to time-varying correlations. Formally, σ -DFC is calculated by first estimating the time-varying covariance using DCC. Under this approach, the data, (y_t) is modeled as a zero-mean stochastic process with auto-regressive covariance:

$$y_t \sim \mathcal{N}(0, \Sigma_t) \quad (12)$$

with time-varying covariance matrix Σ evolving according to the first-order autoregressive model:

$$\Sigma_t = \Omega + A \circ y_{t-1} y_{t-1}^T + B \circ \Sigma_{t-1}. \quad (13)$$

The matrices Ω, A, B are estimated in DCC using maximum-likelihood. We define the σ -DFC matrix as the standard deviation (over time) of the time-varying correlation matrix Q_t :

$$\sigma\text{-DFC} := SD(Q_t) = \sqrt{\frac{\sum_t (Q_t - E_t[Q_t])^2}{T - 1}} \quad (14)$$

with $E_t[Q_t]$ denoting the sample mean over time. To ensure numerical stability, we repeated the DCC algorithm 10 times per case (simulation or true data) and used the median estimated time-series for time-varying correlations. The excursion measure was calculated according to (Zhang et al., 2018). Reliability was computed for each pair of region's DFC statistics using Fisher's ICC of group-demeaned DFC metrics between scanning session (ICC(2,1) in the Shrout and Fleiss convention (Shrout and Fleiss, 1979)). Overall reliabilities collapsed across all regions were calculated using Image Intraclass Correlation (Shou et al., 2013).

2.9. Sensitivity analyses

We conducted sensitivity analyses in Sec. 3.2.1 to test how the

different mechanisms of ground-truth models (e.g. connections vs. decay) influence the estimates of “connectivity” in MINDy and rsFC. We were particularly interested in how each method responded to local heterogeneity (i.e. are MINDy/rsFC estimates of connection strength sensitive to local model parameters: decay and curvature). For each batch of the sensitivity analyses, we first simulated a resampled “individual” multiple times to generate a distribution of trial-to-trial variability (“within-subject”) in elements of MINDy’s weight matrix and the rsFC matrix. We then held the weights of the ground-truth model constant while resampling either the curvature (α) or decay (D) parameters and calculating MINDy weights and rsFC from simulations of the new model. Changes in the estimated connectivity (weights or rsFC) were deemed significant if they occurred with $p < .05$ for the corresponding “within-subject” distribution.

2.10. Statistical analyses

Statistical testing was primarily within-subject between method/condition (e.g. paired t-tests). We used the conservative Bonferroni method for all multiple-comparison corrections. All reported p-values are calculated for two-tailed tests unless indicated otherwise. We use $p \approx 0$ to denote p-values calculated as less than 10^{-20} for which precise numerical estimates may deteriorate.

3. Results

3.1. Overview of results/approach

The Results of the paper are structured as follows. The first section serves to relate MINDy parameter estimates to resting-state Functional Connectivity (and related partial correlation approaches) in terms of differentiating/identifying sources of individual variation. The “ground-truth” models for validation in this first set of analyses are drawn from the empirical distribution of MINDy parameters to ensure that the resultant simulated data is realistic. The second section directly addresses the potential for overfitting by testing whether MINDy models cross-validate and whether parameters are reliable. The third section demonstrates that MINDy parameters have distinct anatomical gradients consistent with previous, theoretical results (Demirtaş et al., 2019, Wang et al., 2019), and highly conserved individual variation (a feature not present in over-fit models). The fourth section demonstrates models’ predictive validity by reproducing individual differences in resting-state dynamics using the empirical models. In the fifth section, we demonstrate that the approach is robust to measurement noise, preprocessing pipelines, and hemodynamic confounds. This section uses three forms of “ground-truth” models. For initially testing robustness to noise and global hemodynamic variability, we again use parameters drawn from the empirical distribution to ensure maximum realism. In subsequent analyses, however, “ground-truth” parameter values are drawn from random hyper-distributions independent of the data and combined with more nuanced hemodynamics. This step tests model performance with more exotic “ground-truths” and prevents circularity. We also consider an additional case in which the simulated fMRI data is generated from randomly-parameterized neural-mass models (operating at the millisecond-scale) to provide insight into the relationship/limitations of MINDy parameter estimates from fMRI and the underlying synaptic connectivity. In the sixth section (Sec. 3.8), we summarize comparisons with Dynamic Causal Modeling which receive fuller treatment in the SI (Sec. 5.8). The final results section directly assesses data-requirements of MINDy and provides a minimum data quantity (>15 min) to prevent over-fitting.

3.2. MINDy retrieves individual differences

3.2.1. MINDy retrieves individualized connectivity

A key goal of our investigation was to determine whether MINDy was

sufficiently sensitive to reveal individual differences in connectivity weights that have become the focus of recent efforts within the rsFC literature (Laumann et al., 2015, Gordon et al., 2017). We tested the model by reconstructing individual differences in connectivity weights of simulated subjects and comparing them against both classical rsFC and the partial correlation matrix. Simulated subjects were generated by permuting MINDy parameter sets across individuals (see methods). We then simulated the resultant model with process noise and hemodynamics to generate realistic BOLD fMRI time series (see methods; Fig. 1 C; SI Fig. 9 B). This provided a ground-truth set of simulated fMRI data, from which we could compute the rsFC/partial correlation matrices for each “subject”, and also determine the fidelity of recovered parameters (i.e., compared against true parameters used to generate the simulated data). To assess the performance of the model estimation procedure, we considered two metrics: the validity of estimated connectivity weight differences between subjects (Fig. 2 B) and the sensitivity of each procedure to different model components (SI Fig. 17 A). These sensitivity analyses reveal whether each approach (rsFC matrix, partial correlation matrix, or model estimation) misclassifies variation in some other model component (e.g. decay rates) as being due to a change in weights. To better assess sensitivity, we generated data after varying only one model component at a time across the simulated subjects: the weight matrix (W), transfer functions (α) or decay rates (D).

Results indicated that MINDy was able to accurately recover the ground-truth weight matrix for each individual (Fig. 2A and B). Thus, the simulated weight changes that differentiated one individual from another were recovered well by the MINDy parameter estimation approach. Moreover, MINDy weight estimates were found to significantly outperform rsFC and partial correlation measures (computed on the simulated timeseries data) in their ability to accurately recover both the ground-truth connectivity matrix of simulated individual subjects, as well as the differences between individuals (Fig. 2 B; SI Table 5). This finding suggests that the modest relation between rsFC and ground-truth connectivity weights is primarily driven by the group-average connectivity as opposed to individual differences. However, rsFC may be disadvantaged in this comparison as it does not typically permit sparseness commensurate with empirical MINDy weights (Fig. 5A and B). Therefore, we used partial correlations as an additional benchmark. While partial correlations quantitatively improved upon rsFC estimates (single-subject: $R = .537 \pm .032$, inter-subject: $R = .392 \pm .027$), performance remained significantly lower than MINDy (single-subject: $\text{paired-}t(33) = 40.51, p \approx 0$, inter-subject: $\text{paired-}t(33) = 23.62, p \approx 0$).

The above analyses were designed to illustrate the additional utility of MINDy in empirical contexts over the most common current approaches (rsFC and partial correlation). For this reason, we generated ground-truths from the empirical distributions to ensure maximal realism. In later analyses (Sec. 3.8), we compare MINDy to a much closer modeling approach (Spectral DCM; (Razi et al., 2017)). We reserve these comparisons for later as they employ a very different approach to generating ground-truth models: seeking to minimize bias and sample over a wide range of potential ground-truth scenarios. The anatomically-detailed models used in the current section are also too large for Spectral DCM to estimate using available computational resources (Sec. 3.8).

3.2.2. MINDy disentangles sources of individual differences

After we established that MINDy outperforms rsFC and partial correlations in retrieving true individual differences in weights, we benchmarked the sensitivity of each approach to other sources of individual variation. Rather than measuring how well each procedure correctly retrieves connectivity, these tests quantify how well each approach selectively measures connectivity as opposed to other sources of variation (see methods). We quantified sensitivity in terms of how often MINDy and rsFC reported that a connection changed in strength between simulated models, when in reality only the curvature or decay terms were altered (SI Fig. 17 A). Results indicate that MINDy correctly detects the sources of individual variation when due to local changes such as decay

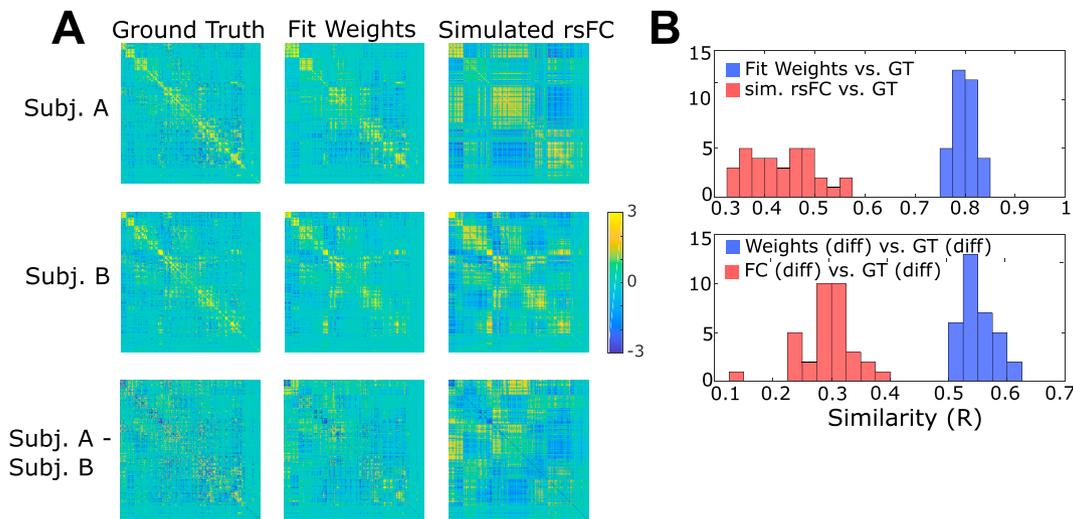


Fig. 2. Ground-truth validation of MINDy and rsFC at the level of single-subject and inter-subject variation. A) First column: Example ground-truth weight matrices for two simulated subjects (top two rows) and the difference between ground-truth weights (bottom-row). Second column: Recovered weight matrices using MINDy for both subjects and their difference. Third column: same as second but using the rsFC. Fit weight matrices and simulated FC matrices are shown in standard-deviation (SD) units with SD computed across the offdiagonal elements of each individual matrix. The ground-truth matrices are displayed in units $2/3 \times SD$ to aid visual comparison. B) Top row: histogram of performance at the simulated single-subject level (correlation with ground-truth [GT]) for MINDy (blue) and rsFC (red). Bottom row: same as top but for predicting the differences in matched-pairs of simulation subjects who differed only in ground-truth connectivity. Simulation subjects were generated by sampling from the distribution of empirical (HCP) MINDy parameters (see Sec. 2.6.1).

rate and transfer function shape, as these have no appreciable impact on MINDy’s connectivity estimates (the false positive rate is near that expected by chance). By contrast, rsFC measurements were highly sensitive to the decay rate ($27.5 \pm 12\%$ of connections changed vs. $7.6 \pm .6\%$ for MINDy, with 5% expected by chance), indicating that some individual differences in FC may be reflective of purely local brain differences as opposed to connectivity between brain regions (SI Fig. 17 A; SI Table 6). These results indicate that MINDy promises to improve both the mechanistic sensitivity and the anatomical accuracy of inferences based upon individual differences in resting-state fMRI. However, it is still the case that resting-state fMRI exhibits *generalized* sensitivity to individual differences in neurobiology, which may suffice for some applications, such as biomarker discovery (see Sec. 3.4).

3.3. MINDy parameters are reliable

In addition to determining the validity of MINDy parameters, it is also critical to establish their empirical reliability. We examined this question

by analyzing measures of test-retest reliability of the parameter estimates obtained for human subjects contributing resting-state scans on two separate days (30 min each). Results indicated that MINDy had high test-retest reliability for all parameter estimates ($> .75$; Fig. 3 A). The reliability of weight estimates was significantly higher than rsFC reliability, although the mean difference was modest ($\Delta R \approx .045$, SI Table 7, SI Table 8). By contrast, the variability in reliability was noticeably smaller for MINDy, meaning that while the mean advantage of MINDy in terms of reliability was modest, its performance was much more consistent across subjects (less variable reliability; SI Table 8).

3.4. MINDy parameters are personalized

For sake of comparison with FC we have thus far emphasized the ability of MINDy to extract brain connectivity. However, MINDy fits brain models, with the connectivity weights (Fig. 5A and B) comprising just one component. For the approach to faithfully reflect the stable differences among individual brains, it is important that it not just accurately

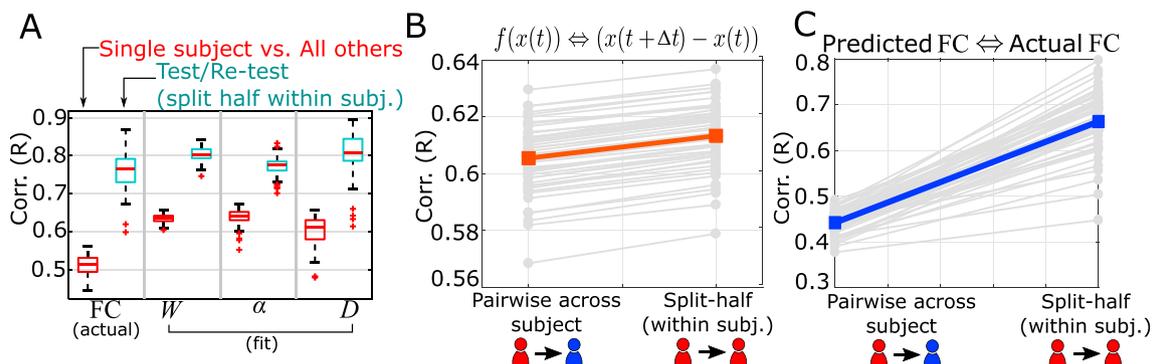


Fig. 3. MINDy parameters and predictions are personalized and reliable. A) Comparison of the test-retest similarity between subjects (red) and the test-retest reliability (blue) for rsFC and the MINDy parameters. B) Goodness-of-fit for a single time-step prediction is uniformly (but minutely) greater for comparing test-retest predictions within a subject vs. between subjects. Performance is in terms of predicting the difference time series. Red line indicates group-mean C) This relationship magnifies across time steps as evidenced by far greater similarity in test-retest predicted FC from model simulations of the same subject vs. different subject. Performance is in terms of predicting the empirical rsFC on a different scanning session. For similarity to the same or both scanning sessions see SI Fig. 16. Blue line indicates mean.

estimates the neural parameters that describe human brains, but that these parameters accurately capture individual differences and predict brain activity. Using the “connectome fingerprinting” approach (Finn et al., 2015), we compared whether MINDy parameter estimates and the combined model uniquely identify individuals within a sample. This analysis was conducted in two ways. First, we computed separate parameter estimates for each individual in each testing day session. Then we examined whether the parameters estimated from one day showed the highest similarity to the same individual on the other day (relative to all other individuals in the dataset; Fig. 3 A). Secondly, we used the estimated model from one day to test whether the estimated parameters provided the best fit to the fMRI data timeseries recorded on the second day, again relative to the estimated parameters from other subjects. Specifically, this second analysis provides a strong form of cross-validation testing and we performed it for both predictions of the empirical timeseries (Fig. 3 B) and for predictions of each subject’s empirical rsFC, both cross-validated across sessions (Fig. 3 C). In all analyses, we found that the best predicting model for every subject was almost always their previously fit model (Table 7). In particular, we achieved 100% accuracy when conducting connectome fingerprinting based on MINDy weight parameters (SI Fig. 17 B), and when computing cross-validated goodness of fit/cross-validated predicted rsFC (Fig. 3B and C). For pairwise analyses of subjects, see SI Fig. (17 F).

Similar patterns emerged but also some important differences, when conducting parallel analyses using rsFC. Replicating prior findings ((Finn et al., 2015)), 100% accuracy was also achieved in connectome fingerprinting (SI Fig. 17E). However, between-subject similarity was significantly lower in the rsFC analysis. Conversely, in rsFC the distinction between across-sessions within-individual similarity scores (i.e. test-retest similarity) and the average similarity obtained between subjects was greater than that observed in the MINDy model weights (SI Table 7). These results suggest that rsFC may actually generate an exaggerated picture of the idiosyncratic nature of connectivity, since MINDy individual differences are partitioned not only into weights, but also into other mechanistic parameters that are attributed locally, to the

node/parcel (i.e., the decay [D] and curvature [α] parameters). In other words, MINDy may provide a richer and more variegated perspective on the nature of individuality, than what can be obtained with rsFC which lumps together what may be multiple dimensions of individual difference, into a simple, undifferentiated measure. For applications such as biomarker discovery, these properties may not be relevant in that the apparent magnification of individual differences in rsFC over MINDy weights could prove beneficial despite the mechanistic ambiguity of rsFC. However, we also note that MINDy provides additional parameters (curvature and decay) which may also prove useful for biomarker discovery. Lastly, the relevant dimensions for biomarker discovery are in terms of separating phenotypes, rather than separating all individuals. Since MINDy can robustly separate individuals, it has the potential to influence biomarker discovery, but whether it possesses quantitative advantages over rsFC will need to be investigated in the context of explicit biomarker questions (and may be phenotype-specific).

3.5. Novel MINDy parameters show reliable individual and anatomic variation

Interestingly, we observed important additional functional utility from examining the novel MINDy parameters that are unavailable in standard rsFC. With regard to individual variation and fingerprinting analyses, we found that even ignoring the weights completely, the transfer function curvature parameter (α) associated with each node showed high consistency across sessions within an individual, and also unique patterns across individuals, such that 100% accuracy could also be achieved in fingerprinting analyses (Fig. 3). A slightly lower accuracy (94.3%) was observed when using the MINDy decay (D) parameters, though even here performance was still significantly above chance (1.89%) in identifying individuals (Fig. 3 A; Table 7). Pair-wise, between-subject, comparisons of similarity in these parameters are reported in SI Fig. (17 B-E).

We followed-up on the identification of reliable individual differences

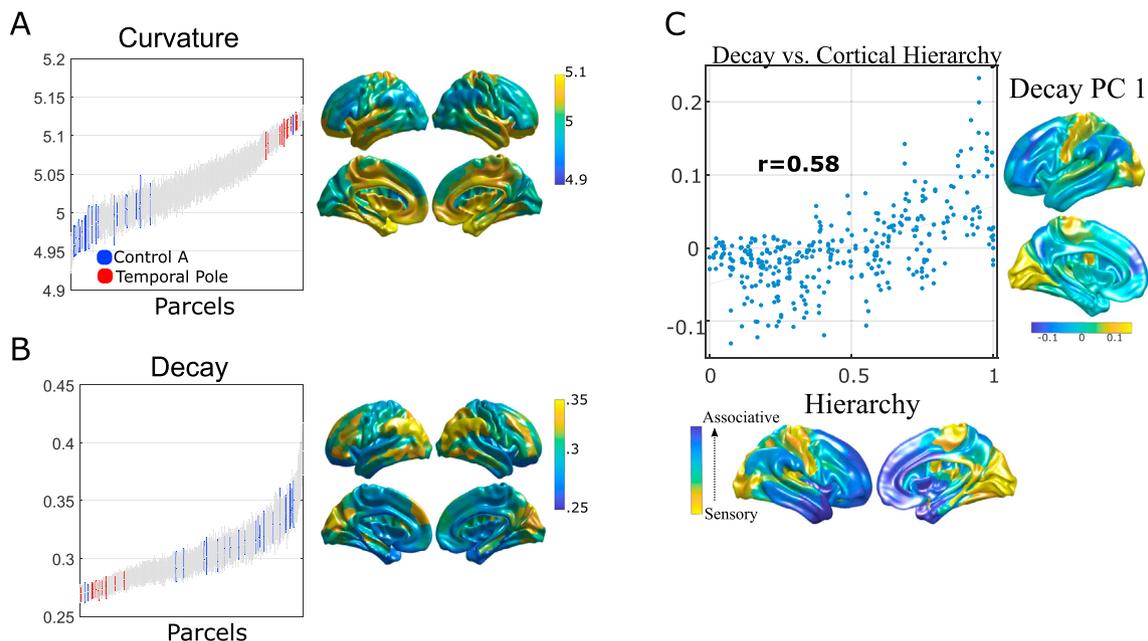


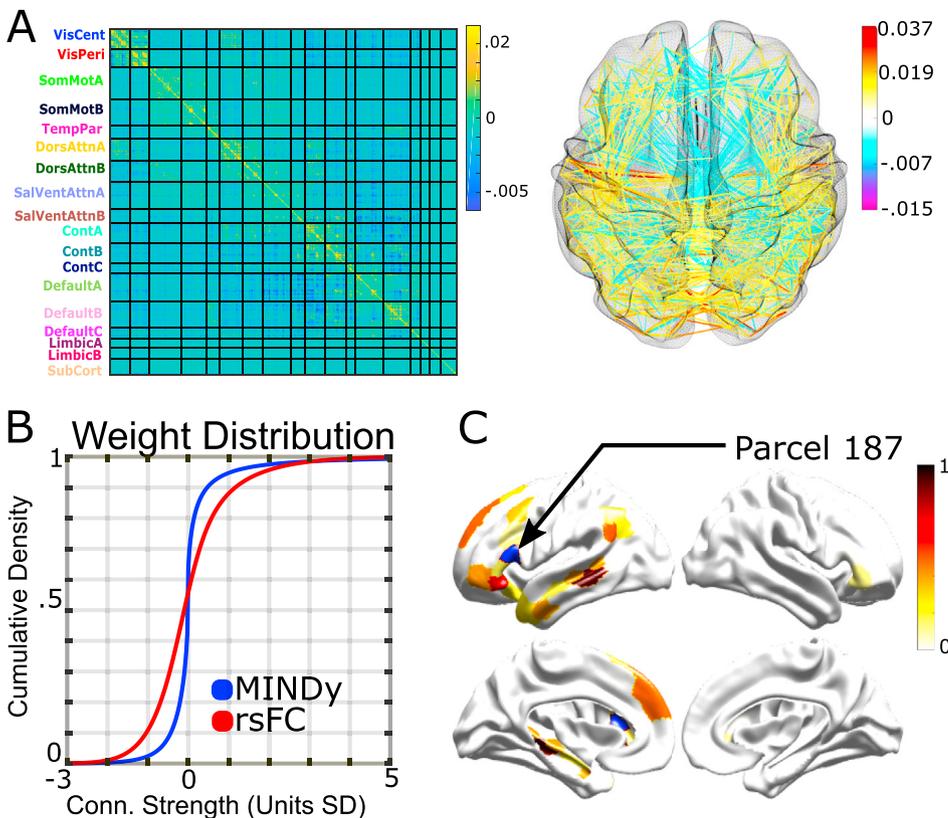
Fig. 4. Local MINDy parameters display consistent anatomical distributions. A) The curvature-parameter displays network structure and is consistent across subjects at the finer parcel level. Parcels are ordered from least to greatest value for the curvature parameter (α) averaged across subjects and scanning sessions. Surface plots are for mean value. Two representative brain networks are highlighted (Control-A in red and Limbic-(Temporal Pole) in blue) to illustrate anatomical gradients in this parameter. B) Same as A but for the decay parameter (D). C) Correlation between the first principal component of MINDy decay and “hierarchical heterogeneity” provided by Demirtaş and colleagues (Demirtaş et al., 2019) based upon erf transform of the T1/T2 ratio (MMP parcelation). This measure has been theorized to reflect a hierarchy of cognitive abstraction from sensory to associative cortices.

through MINDy, by conducting exploratory analyses to examine which brain regions/connections exhibited the greatest inter-individual variability (SI Sec. 5.9). We found that the curvature parameter had greatest relative variability in prefrontal cortex, particularly inferior frontal gyrus (SI Fig. 13A), while the decay parameter had high variability in visual regions, the “hand” portion of post-central gyrus, and medial prefrontal cortex (SI Fig. 13B). Connections within the visual networks had the lowest individual variability while connections to/from the Temporal-Parietal network had the greatest (SI Fig. 13C and D). Although these initial findings are intriguing, due to sample size/bias considerations and the exploratory nature of these analyses, we view them as a launching pad for future insights rather than basic neuroscientific results per se (see SI Sec. 5.9).

Although the above analyses focused on individual differences in the unique MINDy parameters, these parameters also exhibited common patterns across individuals (SI Fig. 11 B,F) that revealed interesting anatomical structure and gradients (Fig. 4 A-C; SI Fig. 11 B,F). These may reflect regional variation in intrinsic dynamics (D) and efferent signaling (α) that vary across brain networks (SI Fig. 11 A,E), but also exhibit consistency even at the finer within-network scale (Fig. 4 A,B; SI Fig. 11 B,F). For example, most nodes within the Temporal-Parietal network showed high curvature, but also low decay parameters; in contrast, in nodes of the Control (A) network, the curvature parameter tended to be low, whereas the decay parameter was high. Group-mean values show the same anatomical gradient across the gwMRF ((Schaefer et al., 2017); SI Fig. 11 D,H) and MMP ((Glasser et al., 2016); SI Fig. 11 C,G) atlases. It is important to note that the decay parameter only describes temporal integration at time scales commensurate with fMRI sampling. Thus, the decay parameter should not be conflated with the time-constant of traditional neural mass models just as the latter is distinct from the membrane time constant of neuronal models. Interestingly, the decay parameter in MINDy appears to reflect components of both temporally-extended signal integration and the time-constant of local sub-second integration. Whereas the mean value of the decay parameter correlates with absolute global brain connectivity (i.e.

the sum of absolute values along a row of the rsFC matrix; $r(377) = .911$, $p \approx 0$) the principal dimensions of individual variation (Fig. 4 C, SI Fig. 11 I and J) recreate the hierarchical organization of primate cortex as derived from the T1/T2 ratio map ($r(358) = .583$, $p \approx 0$; using the MMP Hierarchy map by Demirtaş and colleagues (Demirtaş et al., 2019)). As a caveat, it is worth noting that these statistics do not take into account spatial autocorrelation (which is challenging to model, given the large and irregular shape of parcels), which could have contributed in part to the anatomical gradients we observed. This hierarchy has been the subject of recent studies into its relationship with local excitation/inhibition (Wang et al., 2019, Demirtaş et al., 2019) which is one physiological mechanism we suspect underlies the decay construct (see SI 5.1). This hierarchy also predicts the time-scales of local microcircuits, patterns of gene-expression, myelin density, and function (sensory-processing hierarchy; see (Wang et al., 2019, Demirtaş et al., 2019)).

In addition to the curvature and decay parameters, MINDy also differentiates from rsFC in the structure of the weight matrix (W)/connectivity matrix, both in terms of asymmetry (Fig. 5A,C) and sparseness (Fig. 5A and B). The former is a direct consequence of the dynamical systems model that underlies MINDy, which provides an estimate of effective connectivity. Although regularization generally favors sparse solutions, we found that, even without any regularization, the group average Weight matrix was much sparser than rsFC (SI Sec. 5.7). We provide a simple proof-of-concept to illustrate the potential insights that can be gained from investigating such asymmetries. Specifically, MINDy identified a region of left Inferior Frontal Gyrus (IFG) as the parcel with the greatest asymmetry in positive connections. Specifically, this region showed a positive outward-bias in connectivity with the bias primarily exhibited in its feed-forward positive connections to ipsilateral medial temporal lobe, inferior parietal lobule (IPL), and dorsal/ventrolateral PFC (Fig. 5 C). Excitatory connections of the left IFG with temporal cortex are essential features of the “language network” (e.g. (Friederici and Gierhan, 2013)). Additional results revealing other brain regions showing directionality biases in connectivity are reported in SI (Sec.



5.10). In a later section, we explicitly test the robustness of asymmetry estimates and how they are affected by assumptions regarding hemodynamics and model mismatch.

3.6. MINDy predicts individual brain dynamics

3.6.1. MINDy predicts individual differences in dynamic functional connectivity

We next focused our analyses on the dynamic patterns observed in brain activity, since this has been an area of rapidly expanding research interest within the rsFC literature, termed dynamic functional connectivity or DFC (Hutchison et al., 2013, Allen et al., 2014, Betzel et al., 2016, Preti et al., 2017). Critically, the question of whether MINDy models can predict more slowly fluctuating temporal patterns in the recorded brain data for individual subjects is qualitatively distinct from the ability to predict activity over very short timescales (i.e., 1-step). This is because small biases in fitting individual time points can lead to very different long-term dynamics (e.g. compare panels B and C in Fig. 3, which reflect short and long-term predictions, respectively). To test model accuracy in capturing longer-term dynamic patterns, we used fitted model parameters for each subject to then generate simulated fMRI timeseries, injecting noise at each timestep to create greater variability (see Methods). We then used this simulated timeseries to identify the temporal evolution of short-term correlations between brain regions and compared results with those obtained from the recorded data. Correlation timeseries were estimated using Dynamic Conditional Correlation (DCC; (Engle, 2002)), a method which has been recently shown to improve reliability in the HCP data-set as compared to sliding-window estimates (Choe et al., 2017). We then attempted to recreate DFC measures of individual subjects which have shown the greatest reliability in the actual data. Recent reliability analyses have indicated that simple statistics of temporal variation in individual correlation pairs such as

standard deviation of the conditional correlation time-series (Choe et al., 2017) and excursion (Zhang et al., 2018) are more reliable than state-based descriptions for the HCP resting-state data (Choe et al., 2017). Therefore, we used these measures (see Sec. 2.8 for equations) to validate dynamics within the model. To avoid confusion, we use the term σ -DFC to refer to the temporal standard-deviation of time-varying correlations, which is used as a measure of DFC. Alternatively, the σ -DFC may be conceptualized as the signal power of the moving-correlation time series and has proven to be one of the most reliable measures of DFC (Choe et al., 2017). MINDy performed slightly better on recreating another reliable DFC statistic, group-average excursion, so we chose to be conservative and display the results from σ -DFC for main-text figures rather than using excursion DFC. Results using excursion DFC and the corresponding figures are provided in SI (Fig. 18).

Results indicate that individual differences in the simulated dynamics of models fit to separate test-retest sessions are at least as reliable as summary DFC measures of individual differences in the original data (SI Fig. 18A and B). The image intraclass correlation (I2C2, (Shou et al., 2013)) for the model was 0.555 for σ -DFC and 0.481 for excursion. In the original experimental data, I2C2 reliabilities were 0.527 for σ -DFC and 0.380 for excursion. Moreover, individual differences in the DFC of simulated models were highly correlated with those of the original data for most region-pairs (Fig. 6 A). Lastly, we analyzed whether the simulated data recreates the central tendency of observed data. In general, the group-mean σ -DFC (SI Fig. 18D) and excursion (SI Fig. 18 E) estimates were highly similar between the simulated and observed data for both the σ -DFC (Fig. 6 B; $r(4948) = .761$) and excursion metrics (SI Fig. 18 C; $r(4948) = .836$). Thus, MINDy models recreate measures of DFC at the level of both individual differences and the group-level. Moreover, in some cases (e.g. the excursion metric), MINDy models (fit separately to each session) generate more reliable estimates than those of the original data (SI Fig. 18B). A main advantage of the model in this regard is likely

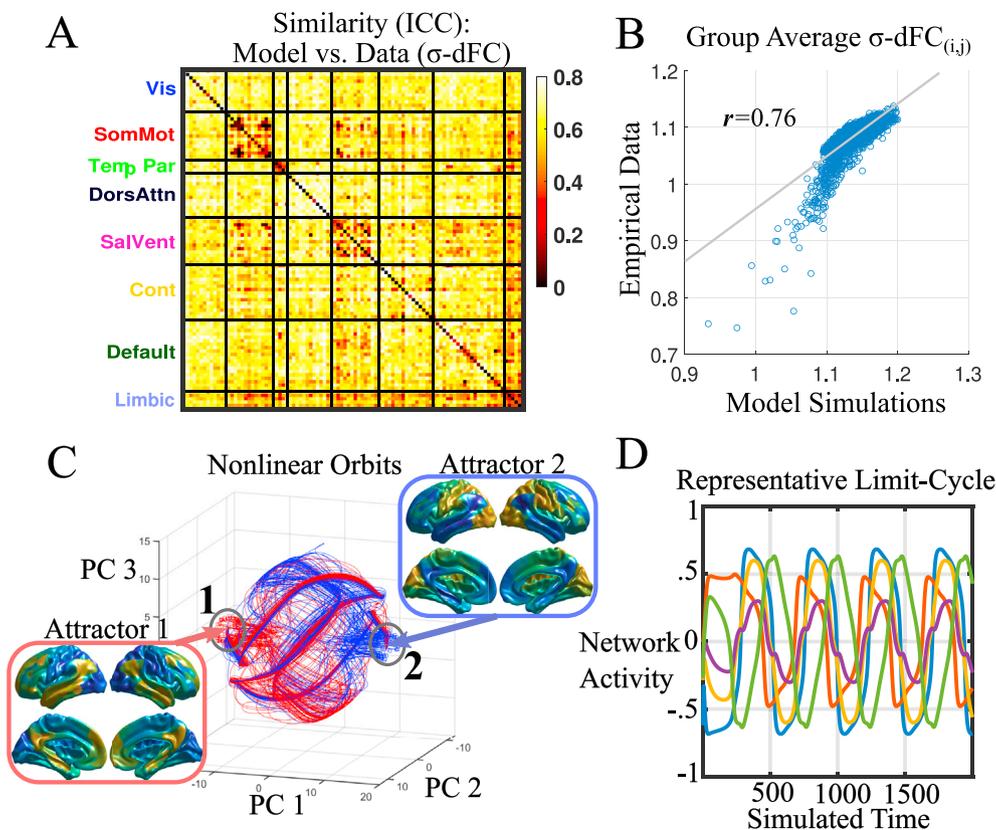


Fig. 6. MINDy models predict individual variation and central tendency of pairwise dynamic functional connectivity (DFC) measures. A) Similarity between model and data for predicting each subject's σ -DFC for each pair of brain regions (using the 100-parcel atlas from (Schaefer et al., 2017) and collapsing the 17-network grouping down to 8). B) Scatterplot and Pearson correlation of group-average σ -DFC for data vs. model. C) Evidence of non-trivial dynamics in MINDy models. Example phase portrait of one subject projected onto the first 3 principal components. Complex orbits link neighborhoods of attractors 1 and 2 (orbits starting from each neighborhood are colored red, blue respectively). Inset figures show these attractors projected onto the brain. D) Example deterministic time-series for a limit cycle in a different representative subject averaged across five networks (Visual [blue], SomMot. [red], Dorsal Attn. [yellow], Saliency [purple], and Control [green]). These deterministic dynamics demonstrate significant nonlinearity but are qualitatively different from the simulated model dynamics (e.g. for computing DFC) which include process noise.

due to the ability to simulate an arbitrarily large amount of data with the model that is also free from nuisance signals/motion.

3.6.2. MINDy models generate non-trivial dynamics

In the previous section we demonstrated that MINDy predicts individual differences in nonstationary dynamics. This finding suggests that the nonlinearities in MINDy are able to account for some features of the data (nonstationarity) that are mathematically absent from linear models. From a dynamics perspective, non-pathological (Schur-stable) linear models predict that spontaneous brain activity consists of noise-driven fluctuations about a single equilibrium. The model parameters for a linear system (e.g. “effective connectivity” in DCM) shape the spatiotemporal statistics of these fluctuations and in the case of white-noise excitation result in a unimodal distribution about the equilibrium in question. Although many nonlinear systems exhibit exotic behavior (e.g. chaos), some systems are dominated by a single equilibrium and may thus possess dynamics that are similar to a linear system. Therefore, we tested whether empirical MINDy models exhibit nontrivial dynamics in the absence of noise (see SI Sec. 5.11). We found that all subjects’ models were dominated by nontrivial dynamics (multistability, homo/heteroclinic cycles, limit cycles, etc.). Example nonlinear dynamics for two representative subjects are provided (Fig. 6C and D), although a thorough characterization of each model’s full phase space is beyond our current scope (see SI Sec. 5.11). Nonetheless, we were able to formally demonstrate that no subject exhibits trivial dynamics (SI Fig. 15A and B; Proposition 2). We conclude that the nonlinearity of MINDy models is not superficial, but rather generates topologically significant dynamics which shape model behavior.

3.7. MINDy is robust

3.7.1. MINDy is robust to measurement noise

We addressed the degree to which MINDy fitted parameters are

influenced by potential sources of contamination or artifact in the observed fMRI data. Resting-state fMRI data is thought to be vulnerable to three main contaminants: noise in the BOLD signal, biases induced from post-processing pipelines that attempt to remove this noise, and idiosyncratic variation in the hemodynamic response function that relates the BOLD signal to underlying brain activity. For the first case, we considered two sources of noise in the BOLD signal: additive measurement error and motion artifact. The former case can result from random fluctuations in magnetic susceptibility, blood flow, and responsiveness of radiofrequency coils among other factors. We examined this issue using the ground-truth simulations described above, but systematically varying the amount of measurement noise added at each time-step. This approach allowed us to determine how strongly these various sources of noise impacted the ability of MINDy to recover the ground-truth parameters. Results indicated that although the performance of MINDy decreased with the amount of noise added (Fig. 7 A), similarity to the ground-truth values generally remained high. Additional levels of noise are plotted in SI Fig. 20. At the highest level of noise considered, Weight and Decay parameters correlated $R \approx 0.7$ with ground-truth, while the curvature parameter correlated $R \approx 0.6$. We note that empirical data exhibiting such a high level of noise would (hopefully) fail quality control.

3.7.2. MINDy is robust to individual differences in motion

We next examined the impact of motion on MINDy estimates. In this case, we used three standard measures of motion that were derived from the observed fMRI timeseries data: 1) the number of total frames censored due to crossing critical values of frame-wise displacement or DVARS (see Methods), 2) the median absolute framewise-displacement of the subject’s head across scanning sessions, and 3) the spatial standard deviation of temporal difference images (DVARS) (Power et al., 2012). We then examined whether variability in these parameters across individuals contributed to the quality of MINDy parameter estimation

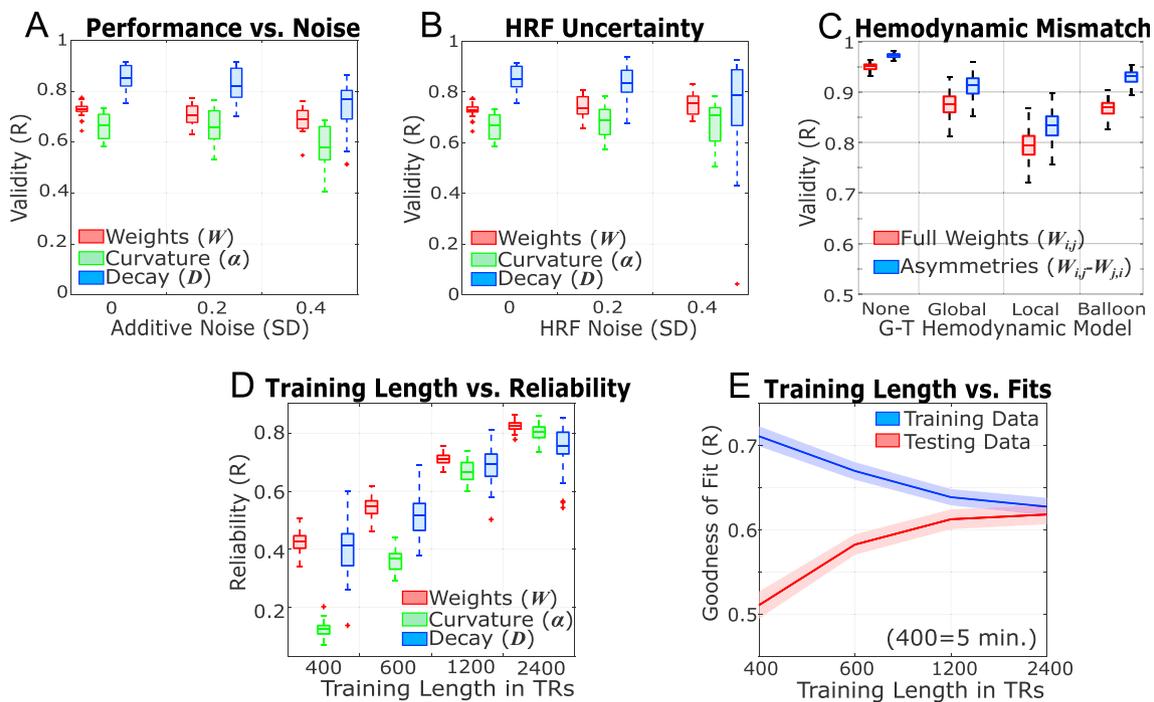


Fig. 7. A) Increasing the amount of additive measurement noise slightly decreases MINDy performance in recovering ground-truth parameters. B) Mean performance is unaffected by the uncertainty in the ground-truth HRF, although performance does become more variable (see SI Fig. 20H). C) MINDy performance in retrieving the ground-truth weight matrix (original matrix: red; asymmetric part ($W_{ij} - W_{ji}$): blue) under mismatch between the ground-truth hemodynamics and the canonical HRF assumed by MINDy (left to right: No hemodynamic modeling, random spatially homogeneous HRF, spatially heterogeneous HRF, nonlinear Balloon-Windkessel model). D) Test-retest reliability of MINDy parameters as the amount of (contiguous) training data is varied. E) MINDy goodness-of-fit for 1-step prediction in training data (blue) and cross-validated with another scanning session (red). The difference between these lines indicates the degree of overfitting. Shading indicates standard deviation.

and individuation, using test-retest reliability (of estimated parameters from each session) as the index of quality. If MINDy estimated parameters reflected vulnerability to the degree of motion present in an individual, then we would expect higher test-retest reliability in the parameters for the individuals with the lowest estimated motion (e.g., highest data quality). Instead, we found that test-retest reliability was relatively un-impacted by any measure of motion (SI Fig. 19B–D, I). A parallel analysis used cross-validated fit, in which MINDy parameters were estimated from one session, and then used to predict data in the held-out session, computing goodness-of-fit of the model to the observed data in this session (in terms of variance explained). In this case, we examined a subset of participants that had relatively low motion in one session, but relatively higher motion in the other compared against a second subset that had similar levels of motion in both sessions. If the increased motion in this latter session was problematic, it should reduce the goodness-of-fit (either when used for parameter estimation or when used for cross-validation in the held-out session). In fact, the cross-validated fits were relatively similar in each group (SI Fig. 19 E,J). Together, these results suggest that participant motion (within a reasonable range) may not be strong factor in determining how well MINDy model parameters can be estimated from observed fMRI data timeseries.

3.7.3. MINDy is robust to pre-processing pipelines

We next examined whether secondary data pre-processing pipelines, which are typically applied to rsFC data prior to analysis, produce biases on MINDy parameter estimates, again examining this issue in terms of test-retest reliability. We considered three variants of a standard published preprocessing pipeline (Siegel et al., 2017), one with motion-correction only, one which adds to this CompCor (a standard method that removes noise components associated with white matter and CSF; (Behzadi et al., 2007)), and a final, full variant that additionally includes global signal regression (GSR; (Aguirre et al., 1998b)). We compared test-retest reliability for data-processed with each pipeline (SI Fig. 21 A) and the similarity of parameter estimates obtained when the same data were processed using different pipelines (SI Fig. 21 B). Results indicated that MINDy parameters had high test-retest reliability regardless of preprocessing choices (all $R > .7$, SI Fig. 21 A) and that similar parameter estimates are obtained regardless of preprocessing choices (all $r > .85$, SI Fig. 21 B). By comparison, when a parallel analysis was conducted on rsFC values, the rsFC parameters showed lower test-retest reliability, particularly when more pre-processing was performed on the data, and showed a larger impact of a change in pre-processing on test-test reliability. A direct comparison of the test-retest of MINDy weight parameters relative to rsFC revealed that these were significantly higher (all p 's $< .05$), were more consistent (lower variance of reliability) across pipelines (all p 's $< .001$; Table 8; SI Fig. 21 A), and were less impacted by changing preprocessing pipelines (all p 's $< .001$; SI Fig. 21 B). Together, this set of analyses indicate that the choice of preprocessing pipeline will not have a large effect on estimated MINDy parameters.

3.7.4. MINDy is robust to global hemodynamics

Lastly, we considered the effect of poor estimation of the hemodynamic response function (HRF). Currently, for simplification, the MINDy estimation procedure assumes a canonical HRF model that is constant across individuals and parcels (although we have recently begun to explore the effect of relaxing this assumption, and estimating a different HRF for each parcel and individual; (Singh et al., 2020)). Many other fMRI models also assume a canonical HRF (e.g. regression-DCM; (Frässle et al., 2017)). However, existing literature suggests that this assumption is likely to be incorrect (Aguirre et al., 1998a, Lin et al., 2018). To examine the impact of mis-fitting the HRF, we modeled a variety of ground-truth scenarios. The first set of ground-truth simulations were randomly parameterized according to the empirical MINDy distribution and activity timeseries were convolved with spatially homogeneous, but randomly parameterized HRFs with incrementally greater variability (SI Fig. 20 D). We then attempted to recover MINDy parameters while again

assuming the fixed, canonical HRF model (Friston et al., 1998). Results of this analysis suggest that, on average, the MINDy parameters recovered from this analysis remain consistently similar to the ground truth parameters (mean similarity of all parameters, R-value \approx 0.75, Fig. 7 B). However, the variability of the fits increased across simulations, as the HRF became more variable across regions and individuals (SI Fig. 20 H).

3.7.5. MINDy parameters are robust to model mismatch

We also considered the effect of violations of the MINDy model in terms of the underlying neural models (MINDy vs. Hopfield, neural mass) and neurovasculature (spatially heterogeneous HRFs and nonlinear hemodynamics). These effects are expected to be most pronounced in estimating asymmetric connections as unaccounted lags can potentially reverse the direction of inferred causality in many other techniques, such as Granger Causality. For the next set of simulations, we generated complex networks from a non-empirical hyperdistribution whose characteristic parameters were randomly sampled at each run. This approach allowed us to sample over a wide range of qualitatively different network structures (Sec. 2.6.2) and these simulations did not depend upon previous empirically-fit MINDy models. We tested the ability of MINDy to recover the weight parameter (Fig. 7 C) from a simple rate model (tanh transfer function) with four levels of hemodynamic variability: 1) no hemodynamics, 2) random, spatially-uniform HRF, 3) random, spatially-heterogeneous HRF, and 4) nonlinear hemodynamics simulated by the Balloon-Windkessel model (Friston et al., 2000, Demirtaş et al., 2019, Sec. 2.6.4). In the last case, the nonlinear hemodynamic transformation varies implicitly and systematically in space due to spatial variation in the firing-rate distribution. Results indicate that MINDy can recover asymmetric connections of ground-truth networks ($W_{ij} - W_{ji}$) for all cases considered, but performance depends upon the degree of HRF complexity (Fig. 7 C; SI Table 9). When no hemodynamics were included in the model (MINDy received the downsampled neural time-series) performance was near-perfect ($r = .949 \pm .009$ overall; $r = .971 \pm .007$ for asymmetries, $n = 1700$). Performance also was high for random, spatially homogeneous HRF's both overall ($r = .874 \pm .024$) and at estimating asymmetries ($r = .910.023, n = 1600$). Spatial heterogeneity of the HRF decreased MINDy performance in recovering overall ground-truth connectivity ($r = .793 \pm .029$; $t(3071.8) = -86.72, p \approx 0$; unequal-variance), but did not differentially impair the estimation of asymmetries ($r = .832 \pm .028$; $t(3057.7) = 11.74, p \approx 1$, 1-tailed, unequal variance).

We also found that MINDy still performed well in recovering asymmetric connectivity when a nonlinear (Balloon-Windkessel) ground-truth hemodynamic model ($r = .865 \pm .022$ overall; $r = .927 \pm .019$ for asymmetries, $n = 2020$) was used to generate simulated fMRI time-series data as compared to when a spatially homogeneous, linear HRF model was used ($t(3073.3) = 23.03, p \approx 0$; unequal variance). Thus, violations of spatial homogeneity in the hemodynamic response appear more relevant to MINDy than violations of hemodynamic linearity. However, performance was still strong in all cases considered (median $r \geq .80$). We also conducted preliminary tests of MINDy's ability to recover synaptic conductances (weights) from the simulated BOLD signal (Balloon-Windkessel) of a biophysically parameterized neural mass model (Hansen et al., 2015) which evolves at a much faster timescale than the fMRI TR. MINDy was generally able to recover connection weights (synaptic conductance in the neural-mass model) for this case as well ($r = .684.039$ overall). However, unlike in the other simulations, performance in recovering asymmetries ($r = .624 \pm .052$) was lower than that of the overall weight matrix ($\text{paired-}t(1399) = -109.172, p \approx 0$). This result indicates that the difference in time-scales between neuronal and BOLD activity is a more relevant constraint on directional inferences than hemodynamic variability. Although these simulations represent but a small subset of possible ground-truth models, they indicate that the directionality of MINDy connectivity estimates remains largely accurate under violations of the assumed spatially homogeneous hemodynamic response.

3.8. Comparing MINDy with Dynamic Causal Modeling

The earlier analyses, in which we compared MINDy and rsFC (Sec. 3.2.1), serve to demonstrate the potential linkages between methods and how MINDy can resolve some ambiguities inherent in rsFC (e.g., directionality). However, these analyses should not be interpreted as stating that MINDy is unambiguously “better” than rsFC as the two approaches represent fundamentally different constructs. The correlation matrix (rsFC) is a statistical quantification whereas MINDy is an approach for estimating a dynamical-systems model and they may have complementary roles for exploring individual differences/biomarker discovery. In order to benchmark MINDy as a model-fitting technique we compared performance with spectral DCM (Razi et al., 2017) in recovering connectivity weights for a variety of simulated ground-truth scenarios. Spectral DCM (spDCM) is a recently developed Dynamic Causal Modeling (DCM) approach for simultaneously estimating linear dynamical systems and (region-specific) hemodynamic kernels from resting-state fMRI (Razi et al., 2017). To be clear, we view the primary contributions of MINDy relative to modeling approaches such as spDCM in terms of its scalability, biological interpretability, and the ability to predict nonstationary resting-state dynamics. However, the question remains whether these advantages come at the expense of accuracy—i.e. whether MINDy is inferior to DCM within the latter’s scope.

We compared performance of MINDy and spDCM across a variety of ground-truth scenarios (see SI Sec. 5.8) to test whether MINDy performs at least as well as spDCM in the lower-dimensional scenarios in which the latter is applicable (i.e., estimating parameters for a small number of nodes or neural masses). These simulations were specifically designed to reduce bias based upon either model’s assumptions (see SI Tab. 1) and considered ground-truths based upon mesoscale Hopfield-style models (SI Fig. 12A) and biophysical neural mass models (SI Fig. 12B). In the former case, we manipulated the degree of spatial variability in the hemodynamic response (SI Fig. 12C). When arbitrary choices were necessary, we chose the option that empirically favored spDCM. Results support that MINDy’s advantages do not come at a cost to accuracy. In all settings considered, MINDy was at least as accurate, on average, as spDCM and significantly (orders of magnitude) faster. We observed that spDCM was more robust than MINDy to spatial variability in the ground-truth HRF (although see extensions in (Singh et al., 2020)), but even under the most extreme cases considered, MINDy was at least as accurate as spDCM (SI Fig. 12). The empirical examination of run-time overwhelmingly favored MINDy (SI Fig. 12D–F). For example, the largest network we tested contained 16 neural masses (SI Fig. 12D) for which MINDy estimation took 3.5s on average vs. 2.7 h for spDCM. We estimate that fitting spDCM models using our chosen parcellation, involving 419 brain regions/nodes (400 of which are cortical (Schaefer et al., 2017)) would take a minimum of 44 years per model (and likely much longer; see SI Sec. 5.8). We conclude that MINDy’s advantages (scalability, dynamics etc.) do not come at the expense of accuracy relative to contemporary approaches.

3.9. MINDy requires 15–20 min of data

In most fMRI experiments scanner-time is a precious resource and particularly so with sensitive populations. While the Human Connectome Project affords a full 60 min of resting-state scan time, this quantity of data may not be a reasonable expectation for other datasets, so we varied the training data length to determine how much data was necessary for MINDy to reliably estimate models. We first evaluated reliability in terms of test-retest on MINDy parameters estimated from separate scanning days using only a subset of the total data for model fitting. As expected, when the length of data used to estimate parameters increased, the test-retest reliability of the estimated parameters also increased, up to the maximum interval considered (30 min). Nevertheless, acceptable levels of reliability ($R \approx .7$) were obtained with 15 min of data (Fig. 7 D). We next examined bias or overfitting of MINDy parameters by comparing the

fit to the trained data, relative to a cross-validation approach, examining the fit to held-out (testing) data. As would be expected with over-fitting bias, as the length of the training data increased, the fit to the trained data decreased. In contrast, the fit to the held-out (test) data increased, and the two values converged at around 15 min of data in training and test sets (Fig. 7 E). Thus, we do believe the current method is too prone to over-fitting biases and unacceptably low reliability with fewer than 15 min of total scan time using the HCP scanning parameters and 419 parcels. However, since the data does not need to be acquired in a single continuous run, we believe that this requirement is reasonable and in concert with current recommendations for rsFC analyses (Laumann et al., 2015, Gordon et al., 2017). Future study with other fMRI acquisition techniques may illuminate how data-requirements change with sampling rate (e.g. shorter TRs may potentially compensate for less scan time).

4. Discussion

4.1. Relationship with functional connectivity

There are three primary advantages to using MINDy over rsFC. First, rsFC is limited as a statistical descriptive model. This means that even though rsFC may be found to be reliable and powerful as a biomarker that can characterize individuals and effects of experimental variables, it is unable to predict how the nature of an experimental manipulation relates to the observed changes in brain activity. By contrast, MINDy is a true mechanistic causal neural model, which attempts to capture the physical processes underlying resting-state brain activity in terms of neurobiologically realistic interactions and nonlinear dynamics (Breakspear, 2017). This feature is powerful, as it enables investigators to perform exploratory analysis in how altering a physical component of the brain (e.g. the connection between two brain regions) will affect brain activity (Jirsa et al., 2017).

Second, MINDy provides a richer description of brain mechanisms than rsFC. While rsFC and MINDy both attempt to parameterize the connection strength between brain regions, MINDy also describes the local mechanisms that govern how each brain region behaves. Neural processes are thought to involve the combination of anatomically local computations and spatially-extended signal propagation, so it is important that descriptions of brain activity be able to define the degree to which this activity is generated by local vs. distributed mechanisms. Although the elements of the rsFC matrix are often interpreted as reflecting interregional components of neural processing, we have demonstrated that the rsFC is also sensitive to purely local characteristics of brain regions, such as their decay rate (SI Fig. 17 A). Conversely, we have demonstrated that both the transfer function and decay rate of brain regions can also serve as reliable markers of individual differences and anatomical structure. By using MINDy, researchers can identify which neural mechanisms (i.e. which of MINDy’s parameters) give rise to individual differences of interest.

Lastly, MINDy greatly improves upon rsFC’s characterizations of effective connectivity between brain regions. Unlike the elements of a correlation matrix, MINDy’s weight parameters can describe asymmetric connectivity strengths and thus make inferences regarding the directional flow of activity between brain regions. We provide tantalizing illustration of the potential utility of these types of findings (e.g., Fig. 5 C, SI Fig. 14). Further, we demonstrated that MINDy may prove a more valid measure of brain connectivity and individual differences in connectivity than rsFC (Fig. 2 E).

4.2. Relationship with other models

There are currently two classes of generative models used to study fMRI: proxy-parameterized neural-mass models (e.g. models using diffusion-imaging data as a proxy for synaptic efficacy Sanz Leon et al., 2013, Jirsa et al., 2017) and directly-parameterized linear models (e.g. Dynamic Causal Modeling (Friston et al., 2003, Razi et al., 2017).

Although nonlinear variants of DCM have been proposed for task-fMRI (Stephan et al., 2008), the techniques used for resting-state fMRI (e.g. Razi et al., 2017) are fundamentally linked with the statistics of linear systems. These two approaches (proxy-parameterized neural-masses and linear models) represent opposite ends of a trade-off between realism and tractability for mesoscale human brain modeling. The first case (proxy-parameterized neural-mass models) excels in terms of interpretability of the model framework, since the state-variables can always be traced back to population firing rates. These models operate at relatively fast time-scales and can produce predictions with the spatial resolution of MRI and the temporal resolution of EEG, which make them a parsimonious and general-purpose investigative tool that can be utilized across temporal scales. Current approaches in this respect are limited, however, in the manner by which these models are parameterized. Even in state-of-the-art techniques (e.g. Demirtaş et al., 2019, Wang et al., 2019) most parameters are fixed a priori (local neural-mass parameters), or determined from diffusion imaging data, with only a limited subset taken from fMRI functional connectivity estimates, typically at the group-average level. Thus, the vast majority of parameters are not sufficiently constrained by the relevant individual-level data, and instead are adapted from measurement of proxy variables, which is likely to limit the accuracy of model predictions. Diffusion imaging data, for instance is inherently unsigned and undirected, so the resultant models are unable to consider hierarchical connection schemes or long-distance connections that depress activity in the post-synaptic targets. Moreover, it remains unknown how to convert units from white matter volume to synaptic conductance even when these assumptions are met. In practice the conversion is performed by choosing a single scaling coefficient, which assumes that this relationship is linear with a universal slope. Due to these sources of uncertainty, proxy-parameterized models do not necessarily fit/predict raw functional time series. To be fair, however, this limitation may not be relevant for all scientific questions (e.g. studying long-term phenomena such as FC, Honey et al., 2009, Demirtaş et al., 2019, Wang et al., 2019).

By contrast, the ability of a model to fit the observed time-series implies that its predictions are valid within the vicinity of observed data. This property holds even when the underlying model is likely to be inaccurate in its long-term predictions. Evidence of this can be seen in the success of dynamic-causal modeling approaches, which can recreate task-driven activity (Friston et al., 2019) despite using a simplified linear model. However, the downside of using a linear modelling framework is that the long-term predictions of these models are most likely inaccurate, given that brain activity in a linear model will always converge to a noise-driven stationary distribution. Thus, even though these models may be more accurate than forward-parameterized neural-mass models in their short-term predictions (by virtue of fitting parameters), the linear form guarantees that they will be unable to capture the extended pattern of brain spatio-temporal dynamics. Analytically, it is known that linear dynamical systems cannot exhibit non-trivial deterministic dynamics and are characterized by a steady-state covariance when driven by noise (which can be calculated by solving a Lyapunov equation). For this reason, they are often employed as surrogate models for testing whether proposed measures of DFC can distinguish between noise-driven trivial (linear) dynamics and those observed in the data (Laumann et al., 2017, Kafashan et al., 2018). Thus, DFC measures which have been shown to be non-spurious through surrogate methods cannot, by definition, be reproduced by a linear dynamical system with or without noise. Likewise, these models will also be limited in their ability to identify the neural mechanisms underlying predictions. Since the model takes a reduced (linear) form, it remains unknown whether the coefficients fit to the linear models are the same as would be retrieved by fitting a more realistic model (e.g. neural mass model). That is not to say that the coefficients are uninterpretable; indeed meaningful predictions have been made by inferring effective connectivity from the model coefficients (e.g. Razi et al., 2017). On the other hand, the models' simplicity may lead to nonlinear components of brain activity being mixed into the linear model

estimates, just as we have shown that intrinsic dynamics influence FC estimates (SI Fig. 17 A).

MINDy attempts to leverage the advantages from both approaches. Like current neural-mass models, MINDy employs a nonlinear dynamical systems model which is capable of generating long-term patterns of brain dynamics. However, MINDy is also a data-driven approach, in that models are fit from the ground-up using functional time-series rather than using surrogate measures such as diffusion imaging (although in principle, such information could be used to initialize or constrain MINDy parameter estimates). From the perspective of biologically-plausible models, MINDy extends parameter fitting from the relatively small number of local parameters that constitute the current state-of-the-art (Demirtaş et al., 2019, Wang et al., 2019) to fitting every parameter in biologically-plausible individualized whole-brain mesoscale models (i.e., increasing the number of fitted parameters by orders of magnitudes). Likewise, MINDy extends data-driven modeling of resting-state data from linear models containing tens of nodes (Razi et al., 2017, Frässle et al., 2017) to nonlinear models containing hundreds. It is also worth noting that the computational innovations made in the optimization process make MINDy parameterization many orders of magnitude faster than comparable techniques (Wang et al., 2019, Razi et al., 2017; see SI Sec. 5.8) despite fitting many more parameters (e.g., over 176,000 free parameters can be estimated in a minute vs. several hours to fit hundreds of parameters). This efficiency has enabled us to interrogate the method's validity and sensitivity in ways that would probably not be computationally feasible for less efficient methods (e.g., building error distributions for sensitivity analyses in Sec. 3.2.1).

4.3. Comparison with diffusion imaging seeded neural mass models

Although we emphasize the ability to generate individualized brain models, previous studies using neural-mass models with weights seeded by diffusion imaging have emphasized predicting group-level data (Honey et al., 2009, Demirtaş et al., 2019, Wang et al., 2019). Two recent studies fit free parameters with the explicit optimization objective of predicting the group-average rsFC matrix (Demirtaş et al., 2019, Wang et al., 2019). By contrast, MINDy seeks to predict the short-term evolution of the neural activity time series for single subjects, which often results in the simulated individual rsFC correlating highly with the empirical rsFC (Fig. 3 C). Averaging across simulated rsFC's produces a group-level simulated rsFC that correlates extremely well with the empirical group-average rsFC ($r(87, 398) = .94$; see SI Fig. 16). As such, the group-average MINDy fit compares very favorably with the analogous measures for diffusion-parameterized models which typically don't surpass $r = .6$ (Honey et al., 2007, Demirtaş et al., 2019, Wang et al., 2019).

4.4. Limitations

There are two primary limitations of MINDy. The first relates to the properties of fMRI data: MINDy is limited by the spatial and temporal resolution at which data is gathered. This means that MINDy is more sensitive to slow interactions that occur over larger spatial scales and is limited to predicting infraslow dynamics (as opposed to higher-frequency bands). Interactions that are more heterogeneous in time or space may also be missed by MINDy as the model assumes that the transfer function is monotone. While the strength of signaling between regions is allowed to vary according to the transfer function, the sign of signaling (inhibitory vs. excitatory) is not. Thus, MINDy cannot describe relationships which, depending upon local activity, change sign from net excitatory to net inhibitory. This feature is inherent in region-based modeling and so this limitation is not unique to MINDy.

A second limitation relates to the model used to specify MINDy. Unlike the conventional neural mass models (Breakspear, 2017), MINDy employs a single population rather than two or more local sub-populations of excitatory and inhibitory neurons. The model does contain local competitive nonlinearities in that the decay term (D) competes with

the recurrent connectivity of W but the precise mechanisms underlying these dynamics are not explicated. By comparison, multipopulation neural mass models contain separate terms for the interactions between subpopulations and the time-constants of firing rate within each subpopulation, both of which likely influence the local parameters of MINDy. Similarly, while MINDy can specify that the directed interaction between a pair of regions is positive, it cannot distinguish between excitatory projections onto an excitatory subpopulation and inhibitory projections onto an inhibitory subpopulation (both of which could be net excitatory; see SI sec 5.1).

4.5. Future applications and directions

We view MINDy models as providing a rich and fertile platform that can be used both for computationally-focused explorations, and as a tool to aid interrogation and analyses of experimental data. The most immediate potential of MINDy is in providing new parameter estimates for studies of individual differences or biomarkers. There is also immediate potential for MINDy in model-driven discovery of resting-state dynamics (e.g. Hansen et al., 2015), in which case MINDy simply replaces diffusion imaging as a method to parameterize neural mass models. The potential benefit of using MINDy over diffusion imaging is that MINDy identifies signed, directed connections in a data-driven manner which may improve realism. Going forward, other applications of MINDy may be in bridging the gap between resting-state characterizations of brain networks and evoked-response models of brain activity during task contexts. We envision two lines of future work in this domain: one in improving estimates of task-evoked effects, and the other concerning the effect of task contexts or cognitive states on brain activity dynamics.

4.5.1. Isolating task-evoked signals

One future use of MINDy may be in improving estimates of task-related brain activity. Current methods of extracting task-related brain signals are based upon comparing BOLD time courses during windows of interest using generalized linear models. However, the effects of task conditions are related to both ongoing brain activity (He, 2013) and intrinsic network structure (Cole et al., 2016). Viewing the brain as a dynamical system, any input to the brain will have downstream consequences, so brain activity observed during task contexts likely contains some mixture of task-evoked activity and its interaction with spontaneous activity. Using MINDy, it may be possible to isolate task-evoked responses by subtracting out what would have been predicted to occur via the resting-state MINDy model. The resultant estimate for task-related activity would be the time-series of MINDy prediction errors (i.e. residuals), ideally adjusted for the model's error at rest. In forthcoming work, we have been using MINDy to estimate task-related activity in this manner, and the initial results strongly indicate that this approach improves the statistical power and temporal specificity of estimated neural events (Wang et al., 2020). Thus, MINDy has the potential to improve estimates of task-evoked activity from fMRI data, although future validation is needed.

4.5.2. Illuminating dynamics

Present results indicate that MINDy is able to replicate some features of infraslow brain-dynamics observed in the data (see Sec. 3.6.1). Although these slower frequency bands have been less studied in task-contexts, growing evidence implicates them in slowly evolving cognitive states such as states of consciousness (Mitra et al., 2015, Mitra et al., 2018) and daydreaming (Kucyi and Davis, 2014). MINDy may benefit future studies in these domains by providing a formal model by which to identify the mechanisms underlying dynamical regimes. Moreover, MINDy may illuminate the behavioral significance of infraslow dynamics. Previous studies have found that timing of pre-cue brain activity and infraslow dynamics interact to predict behavioral performance (Fox et al., 2007, Schölvinck et al., 2012, Sadaghiani et al., 2015), so future characterizations of task-activation may benefit from considering how

exogenous stimuli interact with endogenous neural processes. Generative models of resting-state brain activity may prove critical in these efforts by predicting how endogenous brain states modulate the effects of exogenous perturbations.

5. Conclusion

We have developed a novel and powerful method for constructing whole-brain mesoscale models of individualized brain activity dynamics, termed MINDy. We demonstrate that MINDy models are valid, reliable, and robust, and thus represent an important advance towards the goal of personalized neuroscience. We provide initial illustrations of the potential power and promise of using MINDy models for experimental analysis and computational exploration. It is our hope that other investigators will make use of MINDy models to further test and explore the utility and validity of this approach. Towards that end, we have made MATLAB code and documentation for developing and testing MINDy models available via the primary author's GitHub: <https://github.com/singhmf/MINDy>.

CRedit authorship contribution statement

Matthew F. Singh: Conceptualization, Methodology, Software, Validation, Formal analysis, Writing - original draft, Visualization, Funding acquisition. **Todd S. Braver:** Conceptualization, Writing - original draft, Resources, Funding acquisition, Supervision. **Michael W. Cole:** Conceptualization, Writing - original draft. **ShiNung Ching:** Conceptualization, Writing - original draft, Resources, Funding acquisition, Supervision.

Acknowledgments

MS was funded by NSF-DGE-1143954 from the US National Science Foundation. TB acknowledges R37 MH066078 from the US National Institute of Health. SC holds a Career Award at the Scientific Interface from the Burroughs-Wellcome Fund. Portions of this work were supported by AFOSR 15RT0189, NSF ECCS 1509342 and NSF CMMI 1537015, NSF NCS-FO 1835209 and NIMH Administrative Supplement MH066078-15S1 from the US Air Force Office of Scientific Research, US National Science Foundation, and US National Institute of Mental Health, respectively.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.neuroimage.2020.117046>.

References

- Aguirre, G., Zarahn, E., D'Esposito, M., 1998. The variability of human, bold hemodynamic responses. *Neuroimage* 8 (4), 360–369. <https://doi.org/10.1006/nimg.1998.0369>.
- Aguirre, G., Zarahn, E., D'Esposito, M., 1998. The inferential impact of global signal covariates in functional neuroimaging analyses. *Neuroimage* 8 (3), 302–306. <https://doi.org/10.1006/nimg.1998.0367>.
- Allen, E.A., Damaraju, E., Plis, S.M., Erhard, E.B., Eichele, T., Calhoun, V.D., 2014. Tracking whole-brain connectivity dynamics in the resting state. *Cerebr. Cortex* 24 (3), 663–676. <https://doi.org/10.1093/cercor/bhs352>.
- Ashley, E.A., 2015. The precision medicine initiative: a new national effort. *J. Am. Med. Assoc.* 313 (21), 2119–2120. <https://doi.org/10.1001/jama.2015.3595>.
- Behzadi, Y., Restom, K., Liu, J., Liu, T.T., 2007. A component based noise correction method (compcor) for bold and perfusion based fmri. *Neuroimage* 37 (1), 90–101. <https://doi.org/10.1016/j.neuroimage.2007.04.042>.
- Betzal, R.F., Fukushima, M., He, Y., Zuo, X.-N., Sporns, O., 2016. Dynamic fluctuations coincide with periods of high and low modularity in resting-state functional brain networks. *Neuroimage* 127, 287–297. <https://doi.org/10.1016/j.neuroimage.2015.12.001>.
- Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echoplanar mri. *Magn. Reson. Med.* 34 (4), 537–541. <https://doi.org/10.1002/mrm.1910340409>.
- Breakspear, M., 2017. Dynamic models of large-scale brain activity. *Nat. Neurosci.* 20, 340–352. <https://doi.org/10.1038/nn.4497>.

- Buckner, R.L., Krienen, F.M., Yeo, B.T.T., 2013. Opportunities and limitations of intrinsic functional connectivity mri. *Nat. Neurosci.* 16, 832–837. <https://doi.org/10.1038/nn.3423>.
- Choe, A.S., Nebel, M.B., Barber, A.D., Cohen, J.R., Xu, Y., Pekar, J.J., Caffo, B., Lindquist, M.A., 2017. Comparing test-retest reliability of dynamic functional connectivity methods. *Neuroimage* 158, 155–175. <https://doi.org/10.1016/j.neuroimage.2017.07.005>.
- Cole, M.W., Ito, T., Bassett, D.S., Schultz, D.H., 2016. Activity flow over resting-state networks shapes cognitive task activations. *Nat. Neurosci.* 19 (12), 1718–1726. <https://doi.org/10.1038/nrn.4406>.
- Deco, G., Jirsa, V.K., McIntosh, A.R., 2011. Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nat. Rev. Neurosci.* 12, 43–56. <https://doi.org/10.1038/nrn2961>.
- Demirtaş, M., Burt, J.B., Helmer, M., Ji, J.L., Adkinson, B.D., Glasser, M.F., Van Essen, D.C., Sotiropoulos, S.N., Anticevic, A., Murray, J.D., 2019. Hierarchical heterogeneity across human cortex shapes large-scale neural dynamics. *Neuron* 101 (6), 1181–1194. <https://doi.org/10.1016/j.neuron.2019.01.017>.
- Donoho, D.L., 2006. For most large underdetermined systems of linear equations the minimal l_1 norm solution is also the sparsest solution. *Commun. Pure Appl. Math.* 59 (6), 797–829. <https://doi.org/10.1002/cpa.20132>.
- Dozat, T., 2016. Incorporating nesterov momentum into adam. In: *Proceedings of 4th International Conference on Learning Representations. Workshop Track*.
- Engle, R., 2002. Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *J. Bus. Econ. Stat.* 20 (3), 339–350. <https://doi.org/10.1198/073500102288618487>.
- Essen, D.C.V., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., 2013. The Wu-minn human connectome project: an overview. *Neuroimage* 80, 62–79. <https://doi.org/10.1016/j.neuroimage.2013.05.041>.
- Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X., Constable, R.T., 2015. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nat. Neurosci.* 18, 1664–1671. <https://doi.org/10.1038/nn.4135>.
- Fischl, B., 2012. Freesurfer. *Neuroimage* 62 (2), 774–781. <https://doi.org/10.1016/j.neuroimage.2012.01.021>.
- Fox, M.D., Snyder, A.Z., Vincent, J.L., Raichle, M.E., 2007. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron* 56 (1), 171–184. <https://doi.org/10.1016/j.neuron.2007.08.023>.
- Frässle, S., Lomakina, E.L., Razi, A., Friston, K.J., Buhmann, J.M., Stephan, K.E., 2017. Regression dcm for fmri. *Neuroimage* 155, 406–421.
- Friederici, A.D., Gierhan, S.M., 2013. The language network. *Curr. Opin. Neurobiol.* 23 (2), 250–254. <https://doi.org/10.1016/j.conb.2012.10.002>.
- Friston, K., Fletcher, P., Josephs, O., Holmes, A., Rugg, M., Turner, R., 1998. Event-related fmri: characterizing differential responses. *Neuroimage* 7 (1), 30–40. <https://doi.org/10.1006/nimg.1997.0306>.
- Friston, K., Mechelli, A., Turner, R., Price, C., 2000. Nonlinear responses in fmri: the balloon model, volterra kernels, and other hemodynamics. *Neuroimage* 12 (4), 466–477. <https://doi.org/10.1006/nimg.2000.0630>.
- Friston, K., Harrison, L., Penny, W., 2003. Dynamic causal modelling. *Neuroimage* 19 (4), 1273–1302. [https://doi.org/10.1016/S1053-8119\(03\)00202-7](https://doi.org/10.1016/S1053-8119(03)00202-7).
- Friston, K., Preller, K.H., Mathys, C., Cagnan, H., Heinzle, J., Razi, A., Zeidman, P., 2019. Dynamic causal modelling revisited. *Neuroimage* 199, 730–744. <https://doi.org/10.1016/j.neuroimage.2017.02.045>.
- Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Essen, D.C.V., Jenkinson, M., 2013. The minimal preprocessing pipelines for the human connectome project. *Neuroimage* 80, 105–124. <https://doi.org/10.1016/j.neuroimage.2013.04.127>.
- Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, E., and Yacoub, John, Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Van Essen, D.C., 2016. A multi-modal parcellation of human cerebral cortex. *Nature* 536, 171–178. <https://doi.org/10.1038/nature18933>.
- Gordon, E.M., Laumann, T.O., Gilmore, A.W., Newbold, D.J., Greene, D.J., Berg, J.J., Ortega, M., Hoyt-Drazen, C., Gratton, C., Sun, H., Hampton, J.M., Coalson, R.S., Nguyen, A.L., McDermott, K.B., Shimony, J.S., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., Nelson, S.M., Dosenbach, N.U., 2017. Precision functional mapping of individual human brains. *Neuron* 95 (4), 791–807. <https://doi.org/10.1016/j.neuron.2017.07.011>.
- Griffanti, L., Salimi-Khorshidi, G., Beckmann, C.F., Auerbach, E.J., Douaud, G., Sexton, C.E., Zsoldos, E., Ebmeier, K.P., Filippini, N., Mackay, C.E., Moeller, S., Xu, J., Yacoub, E., Baselli, G., Ugurbil, K., Miller, K.L., Smith, S.M., 2014. Ica-based artefact removal and accelerated fmri acquisition for improved resting state network imaging. *Neuroimage* 95, 232–247. <https://doi.org/10.1016/j.neuroimage.2014.03.034>.
- Hansen, E.C., Battaglia, D., Spiegler, A., Deco, G., Jirsa, V.K., 2015. Functional connectivity dynamics: modeling the switching behavior of the resting state. *Neuroimage* 105, 525–535. <https://doi.org/10.1016/j.neuroimage.2014.11.001>.
- He, B.J., 2013. Spontaneous and task-evoked brain activity negatively interact. *J. Neurosci.* 33 (11), 4672–4682. <https://doi.org/10.1523/JNEUROSCI.2922-12.2013>.
- Holland, P.W., Welsch, R.E., 1977. Robust regression using iteratively reweighted least-squares. *Commun. Stat. Theor. Methods* 6 (9), 813–827. <https://doi.org/10.1080/03610927708827533>.
- Honey, C.J., Kötter, R., Breakspear, M., Sporns, O., 2007. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proc. Natl. Acad. Sci. Unit. States Am.* 104 (24), 10240–10245. <https://doi.org/10.1073/pnas.0701519104>.
- Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., Hagmann, P., 2009. Predicting human resting-state functional connectivity from structural connectivity. *Proc. Natl. Acad. Sci. Unit. States Am.* 106 (6), 2035–2040. <https://doi.org/10.1073/pnas.0811168106>.
- Hopfield, J.J., 1984. Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. Unit. States Am.* 81 (10), 3088–3092. <https://doi.org/10.1073/pnas.81.10.3088>.
- Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, M., Penna, S.D., Duyn, J.H., Glover, G.H., Gonzalez-Castillo, J., Handwerker, D.A., Keilholz, S., Kiviniemi, V., Leopold, D.A., de Pasquale, F., Sporns, O., Walter, M., Chang, C., 2013. Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage* 80, 360–378. <https://doi.org/10.1016/j.neuroimage.2013.05.079>.
- Jirsa, V., Proix, T., Perdikis, D., Woodman, M., Wang, H., Gonzalez-Martinez, J., Bernard, C., Bénar, C., Guye, M., Chauvel, P., Bartolomei, F., 2017. The virtual epileptic patient: individualized whole-brain models of epilepsy spread. *Neuroimage* 145, 377–388. <https://doi.org/10.1016/j.neuroimage.2016.04.049>.
- Kafashan, M., Palanca, B.J.A., Ching, S., 2018. Dimensionality reduction impedes the extraction of dynamic functional connectivity states from fmri recordings of resting wakefulness. *J. Neurosci. Methods* 293, 151–161. <https://doi.org/10.1016/j.jneumeth.2017.09.013>.
- Knock, S., McIntosh, A., Sporns, O., Kötter, R., Hagmann, P., Jirsa, V., 2009. The effects of physiologically plausible connectivity structure on local and global dynamics in large scale brain models. *J. Neurosci. Methods* 3 (1), 86–94.
- Kucyi, A., Davis, K.D., 2014. Dynamic functional connectivity of the default mode network tracks daydreaming. *Neuroimage* 100, 471–480. <https://doi.org/10.1016/j.neuroimage.2014.06.044>.
- Laumann, T.O., Gordon, E.M., Adeyemo, B., Snyder, A.Z., Joo, S.J., Chen, M.-Y., Gilmore, A.W., McDermott, K.B., Nelson, S.M., Dosenbach, N.U., Schlaggar, B.L., Mumford, J.A., Poldrack, R.A., Petersen, S.E., 2015. Functional system and areal organization of a highly sampled individual human brain. *Neuron* 87 (3), 657–670. <https://doi.org/10.1016/j.neuron.2015.06.037>.
- Laumann, T.O., Snyder, A.Z., Mitra, A., Gordon, E.M., Gratton, C., Adeyemo, B., Gilmore, A.W., Nelson, S.M., Berg, J.J., Greene, D.J., McCarthy, J.E., Tagliazucchi, E., Laufs, H., Schlaggar, B.L., Dosenbach, N.U.F., Petersen, S.E., 2017. On the stability of bold fmri correlations. *Cerebr. Cortex* 27 (10), 4719–4732. <https://doi.org/10.1093/cercor/bhw265>.
- Lin, F.-H., Polimeni, J.R., Lin, J.-F.L., Tsai, K.W.-K., Chu, Y.-H., Wu, P.-Y., Li, Y.-T., Hsu, Y.-C., Tsai, S.-Y., Kuo, W.-J., 2018. Relative latency and temporal variability of hemodynamic responses at the human primary visual cortex. *Neuroimage* 164, 194–201. <https://doi.org/10.1016/j.neuroimage.2017.01.041>.
- Lohmann, G., Erfurth, K., Müller, K., Turner, R., 2012. Critical comments on dynamic causal modelling. *Neuroimage* 59 (3), 2322–2329. <https://doi.org/10.1016/j.neuroimage.2011.09.025>.
- Markram, H., 2006. The blue brain project. *Nat. Rev. Neurosci.* 7, 153–160. <https://doi.org/10.1038/nrn1848>.
- Markram, H., Meier, K., Lippert, T., Grillner, S., Frackowiak, R., Dehaene, S., Knoll, A., Sompolinsky, H., Verstreken, K., DeFelipe, J., Grant, S., Changeux, J.-P., Saria, A., 2011. Introducing the human brain project. *Proc. Comput. Sci.* 7, 39–42. <https://doi.org/10.1016/j.procs.2011.12.015>.
- Marreiros, A.C., Daunizeau, J., Kiebel, S.J., Friston, K.J., 2008. Population dynamics: variance and the sigmoid activation function. *Neuroimage* 42 (1), 147–157. <https://doi.org/10.1016/j.neuroimage.2008.04.239>.
- Mastrogioseppe, F., Ostojic, S., 2018. Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron* 99, 609–623. <https://doi.org/10.1016/j.neuron.2018.07.003>.
- Mitra, A., Snyder, A.Z., Tagliazucchi, E., Laufs, H., Raichle, M.E., 2015. Propagated infra-slow intrinsic brain activity reorganizes across wake and slow wave sleep. *eLife* 4, e10781. <https://doi.org/10.7554/eLife.10781>.
- Mitra, A., Kraft, A., Wright, P., Acland, B., Snyder, A.Z., Rosenthal, Z., Czerniewski, L., Bauer, A., Snyder, L., Culver, J., Lee, J.-M., Raichle, M.E., 2018. Spontaneous infra-slow brain activity has unique spatiotemporal dynamics and laminar structure. *Neuron* 98 (2), 297–305. <https://doi.org/10.1016/j.neuron.2018.03.015>.
- Okano, H., Sasaki, E., Yamamori, T., Iriki, A., Shimogori, T., Yamaguchi, Y., Kasai, K., Miyawaki, A., 2016. Brain/minds: a Japanese national brain project for marmoset neuroscience. *Neuron* 92 (3), 582–590. <https://doi.org/10.1016/j.neuron.2016.10.018>.
- Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2012. Spurious but systematic correlations in functional connectivity mri networks arise from subject motion. *Neuroimage* 59 (3), 2142–2154. <https://doi.org/10.1016/j.neuroimage.2011.10.018>.
- Power, J.D., Schlaggar, B.L., Petersen, S.E., 2015. Recent progress and outstanding issues in motion correction in resting state fmri. *Neuroimage* 105, 536–551. <https://doi.org/10.1016/j.neuroimage.2014.10.044>.
- Preti, M.G., Bolton, T.A., Ville, D.V.D., 2017. The dynamic functional connectome: state-of-the-art and perspectives. *Neuroimage* 160, 41–54. <https://doi.org/10.1016/j.neuroimage.2016.12.061>.
- Psaty, B.M., Dekkers, O.M., Cooper, R.S., 2018. Comparison of 2 treatment models: precision medicine and preventive medicine. *J. Am. Med. Assoc.* 320 (8), 751–752. <https://doi.org/10.1001/jama.2018.8377>.
- Razi, A., Seghier, M.L., Zhou, Y., McColgan, P., 2017. Large-scale dcms for resting-state fmri. *Netw. Neurosci.* 1 (3) https://doi.org/10.1162/NETN_a_00015.
- Roebroeck, A., Formisano, E., Goebel, R., 2011. The identification of interacting networks in the brain using fmri: model selection, causality and deconvolution. *Neuroimage* 58 (2), 296–302. <https://doi.org/10.1016/j.neuroimage.2009.09.036>.
- Ryali, S., Supekar, K., Chen, T., Menon, V., 2011. Multivariate dynamical systems models for estimating causal interactions in fmri. *Neuroimage* 54 (2), 807–823. <https://doi.org/10.1016/j.neuroimage.2010.09.052>.

- Sadaghiani, S., Poline, J.-B., Kleinschmidt, A., D'Esposito, M., 2015. Ongoing dynamics in large-scale functional connectivity predict perception. *Proc. Natl. Acad. Sci. Unit. States Am.* 112 (27), 8463–8468. <https://doi.org/10.1073/pnas.1420687112>.
- Salimi-Khorshidi, G., Douaud, G., Beckmann, C.F., Glasser, M.F., Griffanti, L., Smith, S.M., 2014. Automatic denoising of functional mri data: combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage* 90, 449–468. <https://doi.org/10.1016/j.neuroimage.2013.11.046>.
- Sanz Leon, P., Knock, S., Woodman, M., Domide, L., Mersmann, J., McIntosh, A., Jirsa, V., 2013. The virtual brain: a simulator of primate brain network dynamics. *Front. Neuroinf.* 7, 10. <https://doi.org/10.3389/fninf.2013.00010>.
- Satterthwaite, T.D., Xia, C.H., Bassett, D.S., 2018. Personalized neuroscience: common and individual-specific features in functional brain networks. *Neuron* 98 (2), 243–245. <https://doi.org/10.1016/j.neuron.2018.04.007>.
- Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.-N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T., 2017. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebr. Cortex* 1. <https://doi.org/10.1093/cercor/bhx179>.
- Schirner, M., Rothmeier, S., Jirsa, V.K., McIntosh, A.R., Ritter, P., 2015. An automated pipeline for constructing personalized virtual brains from multimodal neuroimaging data. *Neuroimage* 117, 343–357. <https://doi.org/10.1016/j.neuroimage.2015.03.055>.
- Schölvinck, M., Friston, K., Rees, G., 2012. The influence of spontaneous activity on stimulus processing in primary visual cortex. *Neuroimage* 59 (3), 2700–2708. <https://doi.org/10.1016/j.neuroimage.2011.10.066>.
- Shou, H., Eloyan, A., Lee, S., Zipunnikov, V., Crainiceanu, A.N., Nebel, M.B., Caffo, B., Lindquist, M.A., Crainiceanu, C.M., 2013. Quantifying the reliability of image replication studies: the image intraclass correlation coefficient (i2c2). *Cognit. Affect Behav. Neurosci.* 13 (4), 714–724. <https://doi.org/10.3758/s13415-013-0196-0>.
- Shrout, P., Fleiss, J., 1979. Intraclass correlations: uses in assessing rater reliability. *Psychol. Bull.* 86 (2), 420–428.
- Siegel, J.S., Mitra, A., Laumann, T.O., Seitzman, B.A., Raichle, M., Corbetta, M., Snyder, A.Z., 2017. Data quality influences observed links between functional connectivity and behavior. *Cerebr. Cortex* 27 (9), 4492–4502. <https://doi.org/10.1093/cercor/bhw253>.
- Singh, M.F., Wang, A., Braver, T.S., Ching, S., 2020. Scalable surrogate deconvolution for identification of partially-observable systems and brain modeling. *J. Neural. Eng.* <https://doi.org/10.1088/1741-2552/aba07d>. In press.
- Stephan, K.E., Kasper, L., Harrison, L.M., Daunizeau, J., [den Ouden], H.E., Breakspear, M., Friston, K.J., 2008. Nonlinear dynamic causal models for fmri. *Neuroimage* 42 (2), 649–662. <https://doi.org/10.1016/j.neuroimage.2008.04.262>.
- Thomas Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., Fischl, B., Liu, H., Buckner, R.L., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* 106 (3), 1125–1165. <https://doi.org/10.1152/jn.00338.2011>.
- Tu, C., Rocha, R.P., Corbetta, M., Zampieri, S., Zorzi, M., Suweis, S., 2018. Warnings and caveats in brain controllability. *Neuroimage* 176, 83–91. <https://doi.org/10.1016/j.neuroimage.2018.04.010>.
- Wang, P., Kong, R., Kong, X., Liégeois, R., Orban, C., Deco, G., van den Heuvel, M.P., Thomas Yeo, B., 2019. Inversion of a large-scale circuit model reveals a cortical hierarchy in the dynamic resting human brain. *Sci. Adv.* 5 (1) <https://doi.org/10.1126/sciadv.aat7854>.
- Wang, A., Singh, M.F., Etzel, J., Braver, T., 2020. Enhancing task fmri preprocessing via whole-brain neural modeling of intrinsic activity dynamics. In: *Organization for Human Brain Mapping (Virtual)*, p. 2020.
- Weiner, N., 1949. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: with Engineering Applications*. MIT Press, New York City.
- Wilson, H.R., Cowan, J.D., 1972. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.* 12 (1), 1–24. [https://doi.org/10.1016/S0006-3495\(72\)86068-5](https://doi.org/10.1016/S0006-3495(72)86068-5).
- Zhang, C., Baum, S.A., Adduru, V.R., Biswal, B.B., Michael, A.M., 2018. Test-retest reliability of dynamic functional connectivity in resting state fmri. *Neuroimage* 183, 907–918. <https://doi.org/10.1016/j.neuroimage.2018.08.021>.