Enhancing Task fMRI Preprocessing via Individualized Model-Based Filtering of Intrinsic Activity Dynamics

Matthew F. Singh[1,2,3*], Anxu Wang [3], Michael Cole [2], ShiNung Ching[1,4], Todd S. Braver[3]

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

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

3 Department of Psychological and Brain Sciences, Washington University in St. Louis, St. Louis, MO, USA

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

** Corresponding Author. f.singh@wustl.edu 600 S. Euclid Ave., 63110, St. Louis, MO, USA*

Abstract

Brain responses recorded during fMRI are thought to reflect both rapid, stimulus-evoked activity and the propagation of spontaneous activity through brain networks. In the current work we describe a method to improve the estimation of task-evoked brain activity by first "filtering-out" the intrinsic propagation of pre-event activity from the BOLD signal. We do so using Mesoscale Individualized NeuroDynamic (MINDy; [1]) models built from individualized resting-state data (MINDy-based Filtering). After filtering, time-series are analyzed using conventional techniques. Results demonstrate that this simple operation significantly improves the statistical power and temporal precision of estimated group-level effects. Moreover, estimates based upon our technique better generalize between tasks measuring the same construct (cognitive control) and better predict individual differences in behavior. Thus, by subtracting the propagation of previous activity, we obtain better estimates of task-related neural activity.

Keywords: Resting State fMRI, Neural Dynamics, Causal Modeling, Recurrent Neural Networks, Cognitive Control

1. Introduction

Task-related analyses in fMRI typically involve statistical general linear models (GLMs) which seek to identify the amplitude and/or mean timecourse of (BOLD) evoked-responses after removing nuisance covariates. These approaches have proven statistically powerful and characterize much

of the current literature regarding task-induced activation in group-level fMRI analyses. However, over the past two decades, improvements in fMRI acquisitions and the rise of resting-state connectomics ([2]) have given rise to a new literature concerning variability within brain activation across trials, individuals, and/or contexts. Understanding such variability is key to precision neuroscience initiatives, as these studies have the potential to uncover new neural mechanisms and generate stronger brain-behavior linkages at the level of individuals ([3], [4], [5]).

Previous studies in this domain have generated two key findings relevant to the current study: 1) individual differences in intrinsic brain networks predict corresponding differences in BOLD responses ([6], [7], [8], [9]) and 2) the BOLD signal elicited by a stimulus is dependent upon the previous pattern of brain activity ([10]), including spontaneous fluctuations ([11]). We use the term "brain activity" in the latter case to indicate that this history dependence is thought to be neural, rather than solely reflecting potential nonlinearity in the hemodynamic coupling. The first set of findings indicate that inter-subject variability in brain responses may be due to the "flow" ([8]) of evoked activity through subject-specific connectomes. The second set of findings suggest that evoked responses are history-dependent (i.e. reflects underlying dynamics). Thus, the neural activity associated with BOLD is increasingly considered as a nonlinear dynamical system—one in which the spatiotemporal response to an input depends upon its current state, and further, is determined by a set of rules that dictate its temporal evolution ([12]). These dynamical "rules" are a function of subject-specific connectivity and the specific properties local to each brain region ([13], [14]). The manifestation of these dynamics (i.e. trial-to-trial variability in BOLD) are thought to be neural and cognitively-relevant as they predict within-subject behavioral variation ([15]).

This framework contrasts both with current statistical approaches, which treat the neural activity as a noisy autoregressive signal (most GLMs), and with Dynamic Causal Modeling (DCM) approaches, which treat the brain as a linear system (although see [16]). In the current work, we propose a new technique for modeling intrinsic brain dynamics and their contribution to task-evoked activation patterns. This approach leverages MINDy models ([1]) fit to resting-state data for each subject. These models are akin to an abstracted neural mass model containing hundreds of different regions (parcels) spanning the whole brain. Regions interact nonlinearly via a signed, directed connectivity matrix and integrate inputs over time (i.e. form a nonlinear dynamical system). The BOLD signal is modeled via region-specific hemodynamic models, and all parameters (neural and hemodynamic) are directly estimated from each subject's resting-state

scans (a process which takes 1-3 minutes). In prior work ([1], [17]), we have established that MINDy models/parameters are robust, reliable, and predictive ([1]). In the current work, we use these models to estimate intrinsic brain dynamics (i.e. predictions based upon resting-state MINDy models) and subtract them from the observed BOLD, a process which we term MINDy-based Filtering. This procedure more accurately identifies individual differences, and enhances the temporal precision and statistical power in identifying task effects. We also obtain stronger brain-behavior linkages and more generalizable effects across tasks tapping a common cognitive construct (cognitive control demand).

1.1. Filtering Intrinsic Dynamics

The current approach rests upon the ability to model the flow of neural activity between brain areas, as identified via models fit to resting-state brain activity. However, rather than seeking to describe the flow of task-related neural activity (e.g. [8]), our approach acts to censor, or computationally estimate and remove, the flow of task-unrelated (pre-event) activity. To be clear, we perform this operation at every time point and use the whole timeseries for analyses. No information regarding task timing is used in our filter (Fig. 1). However we use the notion of "events" to provide an intuitive motivation for our approach (conversely each timepoint could be considered an "event"). Likewise, our approach does not require an event-related design (see SI 7.3 for block-related analyses). At each time point, the measured neural activity is considered a combination of task-effects manifest over fast time scales and the propagation of brain activity emerging at previous time points. By subtracting the modeled propagation of previously-triggered (e.g. pre-event) activity, we aim to better isolate the influence of each event (time-point).

Our approach is conceptually-similar to a previous study by Fox and colleagues ([11],[18]) which suggested that estimated task-effects could be improved by subtracting spontaneous activity. They demonstrated this possibility in a motor task by subtracting the recorded BOLD in contralateral motor cortex from the task-implicated motor hemisphere. However, the Fox et al. approach ([11],[18]) has not been applied more broadly, since it requires identifying region pairs which are strongly correlated at rest, but only one of which is recruited during task. This dissociation is key as it enabled Fox and colleagues ([11]) to measure intrinsic brain activity (via the contralateral cortex) separately from task-evoked activity in the other hemisphere. By contrast, the current literature overwhelmingly suggests that, for most brain regions and networks, coactivation during resting-state fMRI predicts coactivation during task (e.g. [6], [8], [7]).

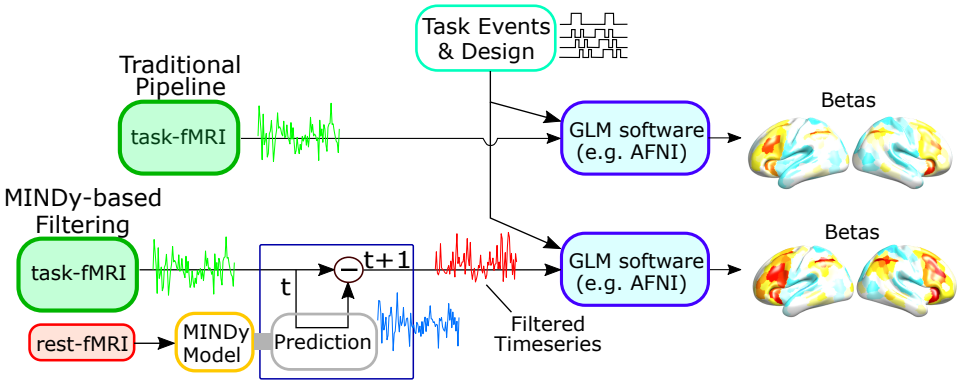

Fig 1. MINDy-based Filtering analysis. MINDy models are estimated from each subject's resting-state data. The MINDy filter makes 1-step (i.e. predicting BOLD at $t+1$ using information at time t) and subtracts the predicted values from the measured BOLD. No information regarding task timing is used in this step. The filtered timeseries are then analyzed using conventional methods.

By contrast, we propose to filter out the intrinsic component of brain activity using model-based predictions. We predict brain activation at each time-point by applying resting-state MINDy models ([1],[17]) to the previous time-step (i.e. 1-step forward predictions) and subtract these predictions to better identify task-evoked changes. Thus, we better isolate event-related brain changes by filtering out the propagation of pre-event activity. As mentioned previously, we use the notion of task "events" to provide an intuitive understanding of why our approach improves fMRI analyses. Our filter does not utilize any prior information regarding task structure (events) and is compatible with any task design (not just event-related designs; see Fig. 1).

1.2. Previous Approaches using DCM

Dynamic Causal Modeling (DCM), by contrast, incorporates the temporal evolution of brain activity and thus can consider the propagation of neural activity through brain networks. Each DCM contains an effective connectivity matrix and a set of extrinsic inputs that describe how task events impinge upon each node of the network ([19]). Many implementations also contain region-specific hemodynamic models and/or an interaction between task events and effective connectivity (i.e., the effective connectivity is parameterized by task events). Although the original DCM models were strongly limited in size, modern implementations ([20], [21]) can consider a much larger number of brain regions (although the computation cost still remains considerable; [20], [1]). However, the DCM methodology also presents several constraints which limit its application. Estimating a DCM model requires pre-specifying the time-series of task effects. This assumption precludes analyses which explore the temporal dynamics of task effects such as Finite Impulse Response (FIR) modeling or nuanced task GLMs,

such as those featuring nuisance regressors (e.g. motion). In addition, all DCM implementations that support whole-brain models (i.e., more than a few regions; [20]) are dependent upon the assumption of stationary linear dynamics ([1]).

1.3. Filtering Instead of Parameterizing

In the current work, we aim to strike a balance between the mechanistic inferences made by DCM and the flexibility of standard analysis techniques. To do so, we generate dynamical systems models of the brain and neurovasculature (as is done in DCM). However, our approach differs substantially from DCM in how we build and utilize these models. Instead of fitting models of the brain and tasks, we propose to fit dynamic models to independent resting-state data for each subject. We then use these models to generate a mathematical filter for each subject that removes, or "partials out", the effects of intrinsic dynamics from BOLD timeseries. The approach uses no information regarding task events and thus functions as a preprocessing step, as opposed to explicitly modeling task events. This feature is advantageous, as the proposed techniques can be inserted into any data preprocessing pipeline with minimal effort, provided that sufficient amount of resting state data (e.g. >15 minutes [1]) has been collected to build MINDy models.

2. Approach

2.1. Model Derivation

Our approach leverages individualized resting-state models in order to estimate task-evoked brain effects, while making minimal modeling assumptions about the underlying mechanisms. We model brain activity (x_t) as a dynamical system containing two components: an intrinsic dynamical component $f(x)$ which is estimated from resting-state models, and an exogenous input component I_t.

$$x_{t+1} = f(x_t) + I_t. \tag{1}$$

The latter component is exogenous with respect to the resting-state model and should not be interpreted as "exogenous to the brain". Rather, I_t represents additional input to each brain region beyond that which is created through intrinsic (resting state) dynamics embedded in $f(x)$. In principle, this technique is compatible with any resting-state model ($f(x_t)$). For the current work, we chose MINDy ([1], [17]) as it is highly scalable, nonlinear, and robust to many nuisance factors. The aim of the current work is to estimate the input (I_t) for task data and to investigate exogenous input as a marker for cognitive states. We do not assume a specific mechanism underlying this input (e.g. recurrent input, inter-regional signaling, neuronal

"noise", or sensory afferents are all possible sources) or any spatial/temporal properties of I_t. Thus, we treat I_t as a latent signal to be estimated (i.e. filtering I_t from BOLD). By contrast, other methods, such as DCM ([19],[22]) assume a time course of I_t (the temporal aspects of I_t) based upon task design and only estimate its relative contribution to each brain area. For this reason, we term our objective MINDy-based Filtering. Although the mechanisms of interest (I_t) are modeled as neural, fMRI measures the hemodynamic BOLD contrast. For this reason, we simultaneously model neural dynamics and the hemodynamics which link neural events to fMRI measurements. We assume that BOLD signal reflects the convolution (denoted "*") of latent neural activity (x_t) with a region-specific Hemodynamic Response Function (HRF; denoted $h_i(t)$) and we estimate the HRF kernels from resting state data ([17]). Thus, for each brain region (parcel "i") our model is:

$$BOLD_t^{(i)} = [h_i * (x_\tau^{(i)} + \eta_\tau^{(i)})]_t + \nu_t \tag{2}$$

We consider noise at the level of the neurovascular coupling η_t and at the level of BOLD measurements (ν_t). These terms are modeled as normal random variables which are independently and identically distributed (iid) between brain regions and time points. Process noise (physiological stochasticity) is not explicitly modeled at the neural level in Eq. 1, as it is absorbed in the unknown inputs I_t. Substituting for x_t (from Eq. 1) and rearranging yields:

$$BOLD_{t+1}^{(i)} - [h_i * f^{(i)}(x)]_t = [h_i * I_\tau^{(i)}]_t + [h_i * \eta_\tau^{(i)}]_t + \nu_t^{(i)}. \tag{3}$$

Thus, the HRF-convolved input $[h * I]_t$ is equal to the difference between measured and predicted BOLD plus additional autocorrelated noise terms. For all current analyses we consider brain states estimated with HRF-convolved estimates of input ($[h * I]_t$) as opposed to the estimates of I_t alone. This step enables the same statistical pipelines (i.e. GLM structure) to analyze original fMRI BOLD data and the HRF-convolved input. As a result, the estimation of $[h * I]_t$ serves as an additional "preprocessing" (filtering) step that can be added to any fMRI pipeline with minimal effort. No information regarding task events is used in estimating I_t, so the same statistical frameworks are applied to model-filtered and original data.

2.2. MINDy-based Filtering

In the current approach, we do not explicitly model different forms of noise. The only noise factor we consider is the measurement noise power in inverting BOLD onto neural activity. Since neurovasculature noise is removed (η_t=0), Wiener deconvolution ([23]) generates the least-mean-square estimate for x_t. The resultant approximation for BOLD-convolved input ($[h * I]_t$) is:

$$[h * I_\tau]_t \approx BOLD_{t+1} - [h * f(h\hat{*}^{-1} BOLD]_\tau]_t \tag{4}$$

With $h\hat{*}^{-1}BOLD$ denoting the Wiener deconvolution of each region's BOLD signal with respect to the corresponding HRF model. Thus, we estimate neural activity by deconvolving BOLD with the region-specific HRF's identified at rest. Predictions are made in terms of neural activity and then re-convolved to produce predictions in terms of BOLD. The difference between measured and predicted BOLD approximates the HRF-convolved input.

3. Methods

3.1. Subjects

Data consisted of fMRI task and resting-state scans for 53 healthy young-adult subjects collected as part of the Dual Mechanisms of Cognitive Control (DMCC) study ([24]). The DMCC participant pool contains a large number of monozygotic and dizygotic twin pairs. However, for these analyses, these characteristics are ignored.

3.2. Scanning Protocol

Each participant took part in three separate scanning sessions which occurred on different days, but all had the same general procedure. Each day, participants provided two resting-state scans of 5 minutes each as well as two scans each for four cognitive tasks: the AX-Continuous Performance Task (AX-CPT), Sternberg Task, Stroop Task, and Cued Task-Switching (Cued-TS). The two scans per task were performed sequentially for each task whereas the two resting-state scans were separated in time (one at the session start and one at end). Each of the task scans (2 per task per day) contained three task-blocks separated by inter-block intervals and lasted approximately 12 minutes. For resting state and task, the two scans per day were split between anterior-posterior and posterior-anterior phase-encoding directions. Scans were performed at 3T with 1.2s TR (multi-band $\times 4$; see [24] for additional details).

3.3. Task Descriptions

We briefly describe the general structure of each of the four cognitive tasks in the "baseline" format which was administered on the first scanning day (see [24] for more details on task design and rationale). Subtle changes to task structure were made on the two following days (subsequent section) but were not relevant to our analyses. The **AX-CPT** task ([25]) involves repeated sequences of cue-probe pairs, in which the response to the probe item is constrained by the preceding contextual cue. Thus, the A-X cue-probe pairing requires a target response and is frequent, leading to strong associations between the cue and probe. However, both the B-X pairing (where "B" refers to any non-X cue) and A-Y pairing (where "Y" refers to any non-X probe) require non target responses. In the **Sternberg**

task ([26]), participants are sequentially presented with short list of words to memorize for that trial (called the memory set; appearing across two encoding screens). After a short retention delay, they are presented with a probe word and must determine if the probe was present in that trial's memory set. On some trials, the probe item is termed a "recent negative", in that was not present in the current trial memory set but was present in the memory set from the preceding trial. In the current implementation of the **Stroop task**, subjects are asked to verbally report the font color in which probes are displayed ([27]). Each probe is itself a color-word, and can either be congruent (font color is the same as the color word, e.g., BLUE in blue font) or incongruent (font color is different from the color-word name; e.g., BLUE in red font). Lastly, during **Cued Task-Switching** (Cued-TS, [28]) participants are pre-cued to attend to either the number or letter component of a subsequent probe (combined letter + digit). In "attend-number" trials, participants indicate whether the digital component of a probe is even vs. odd. In "attend-letter" trials, participants indicate whether the letter component is a consonant vs. vowel. The probe can be either congruent (both letter and digit are associated with the same response) or incongruent (the letter and digit are associated with different responses). With the exception of the Stroop task, participants report responses using button presses.

3.4. Cognitive Control Demand

The current set of trial-based analyses center upon the ability to identify neural signatures of cognitive control. Although cognitive control is a heterogeneous construct, we specifically studied the conflict resolution aspects of cognitive control, so we use the terms control-demand and conflict interchangeably when referring to these tasks, and contrasts between trial types. In particular, we operationally identify cognitive control demand as the difference in neural activity measures during high and low-conflict trials for each task. In the AX-CPT, we contrast BX trials (high conflict) vs. BY (low conflict). The BX trials are high conflict because of the target-association with the X-probe, which require contextual cue information to override. For the Sternberg task, we contrast trials with recent negative probes (high conflict) and trials containing novel negative probes (low-conflict). Thus, recent negative trials are high conflict because the familiarity of the probe, requires information actively maintained in memory to override. In the Stroop task, we contrast incongruent (high conflict) and congruent (low conflict) trials. The incongruent trials are high conflict because the task goals (name the font color) are required to override the dominant tendency to read the color-name. Lastly, in the Cued-TS we also contrast incongruent (high conflict) and congruent (low conflict) trials. The incongruent trials are high conflict because it is critical to process the task cue, in order to

know what response to make (for congruent trials, the same response would be made regardless of the task being performed).

3.5. Task Manipulations

The four tasks (AX-CPT, Sternberg, Stroop, and Cued-TS) were chosen to measure/engage cognitive control. On the first scanning day, participants performed a "baseline" version of each task. On the subsequent days, however, participants performed modified version of each task, meant to promote either proactive or reactive cognitive control strategies. On the two subsequent scans participants performed all the reactive-mode conditions of the tasks on one day and all the proactive-mode conditions of the tasks on another, with the order of proactive vs. reactive days counter-balanced across subjects. In the current work we do not consider the influence of cognitive-control mode and combine data for each task across scanning sessions, to increase statistical power.

3.6. Behavioral Measures

In each task we recorded two behavioral measures: reaction time (RT) and accuracy. Reaction times for button presses were recorded digitally, whereas reaction time for the Stroop task was defined by the duration of silence (time until participant begins a verbal response; see [24]). For the current work, we focused upon the difference in performance measures between trial-types with high cognitive control demand and those with low cognitive control demand (see below). As in previous work with these tasks, we observed lower performance (higher RTs and lower accuracy) on the high demand trials indicative of a cognitive control effect ([24]). For the RT data, we defined cognitive control effects as the difference in normalized RTs between high and low-control trials:

$$RT_{HL} = z(RT_{High}) - z(RT_{Low}) \tag{5}$$

with z denoting z-score normalization. We separately normalized the high and low RT conditions to account for potential heterogeneity of variance between conditions. However, we could not separately normalize accuracy by condition as some of the low-control distributions were near-degenerate (e.g. in one Stroop session over 90% of subjects had 100% accuracy for low-control trials). As such, we normalized accuracy after subtracting high and low-control conditions since the low-control variance was not stable:

$$Acc_{HL} = z(Acc_{High} - Acc_{Low}) \tag{6}$$

As with neural data, we averaged the normalized response times between sessions for each task. Interestingly we found that, unlike RTs, neural data using conventional techniques only predicted errors in the baseline session. Therefore, we only used the baseline error rates for benchmarking (averaged over tasks) and similarly for neural data.

3.7. Pre-processing and Parcellation

Raw resting-state and task data were preprocessed using the same pipeline, implemented with fMRI-prep software ([29],[30]). The whole-brain surface data were then parcellated into 400 cortical parcels defined by the 400 parcel Schaefer atlas (Schaefer [31]; 7-network version). Subcortical volumetric data was divided into 19 regions derived from FreeSurfer ([32]). Motion time-series consisted of the 3-dimensional coordinate changes for rigid-body (brain) rotation and translation (6 total). Motion and linear drift were regressed out of pre-processed resting-state data before MINDy model fitting and from task data prior to filtering. Since motion time-series are also covariates within our task GLMs (as is common), this step does not bias results, as motion is implicitly removed from the unmodeled data during GLM estimation (see below). However, we also implemented controls (see Sec. 3.10) which used this same data (i.e. motion pre-regressed) with conventional analyses.

3.8. Task GLM Analyses

Statistical models of task fMRI were estimated using general linear models (GLM) as implemented in AFNI. The same analyses were performed for both the original task data and the model-subtracted data. The GLM design consisted of a mixed block/event-related design in which trial-type effects were modeled using a modified Finite-Impulse-Response (FIR,[33], [34], [35]) framework (AFNI TENT; [36]), whereas block effects (task vs. inter-block interval) were modeled using a canonical HRF convolved with the block regressors. The TENT bases consisted of overlapping linear-interpolation splines spanning two TRs each with 1TR spacing between bases (i.e. 1TR resolution with stimulus-alignment between TRs). The FIR models were generated by projecting TENT coefficients by the mean TENT basis-set for each trial-type (within-subject). The GLM design also included block onset/offset (modeled with a canonical HRF) and the six motion regressors corresponding to rigid body translation and rotation (3 each). Timepoints containing excessive motion (Framewise Displacement > 0.9mm) were censored from task GLMs. Estimation was performed using the built-in AFNI function "3dREMLfit".

3.9. MINDy Modeling

Mesoscale Individualized NeuroDynamic (MINDy, [1][17]) models were generated from each subject using the parcellated, pre-processed resting-state data for each subject, combined across scanning sessions. Thus, a single MINDy model was estimated for each subject and was used in analyzing task-data across scanning sessions. We simultaneously estimated the neurovascular coupling/HRF and latent brain networks by combining the original MINDy model with Surrogate Deconvolution as in [17]. This

combination simultaneously estimates HRF kernel parameters for each brain region as well as the connectivity matrix, region-specific transfer function shape, and local decay parameter (time-constant). Previous work indicates that the inclusion of Surrogate Deconvolution renders MINDy estimates robust to spatial variation in the HRF. Moreover, the spatial distribution of estimated HRF properties such as time-to-peak are consistent with empirical literature at the group level and are also reliable at the level of individual differences ([17]). Hyperparameters used in MINDy model fitting were identical to previous studies ([1]).

3.10. Control Pipelines

In addition to comparing the proposed pipeline with conventional analyses, we also repeated all task analyses for several control pipelines. These control pipelines considered two factors that might explain results: 1) pre-processing and 2) mechanistic components of the model (SI Sec. 7.2). The MINDy modeling framework assumes that nuisance covariates such as motion and drift have already been removed from time-series prior to model fitting. Therefore, to address #1, we implemented a control in which standard GLM analyses were computed directly upon the fMRI BOLD task timeseries, with motion covariates already regressed out first. The same regressors also appear in the task GLM model (which is shared across all pipelines), but regressing these factors out first will rescale estimated beta-coefficients due to the input normalization performed by many fMRI processing packages (e.g. AFNI). This control ensured that improvements in group-level sensitivity were due to increased similarity of estimated spatiotemporal patterns rather than theoretically uninteresting factors due to pre-processing pipelines. We refer to this control as "pre-regressed" (pre-Reg).

In the SI (Sec. 7.2), we address #2 by considering the influence of anatomically local dynamics vs. interactions between brain regions. This contrast is significant for three reasons. First, it is theoretically significant to distinguish between purely local neural dynamics and inter-regional brain dynamics. Secondly, long distance interactions between brain regions cannot be explained solely in terms of neurovasculature since the regions involved may share anatomically distinct blood supply (i.e. different cerebral arteries). As a result, improvements identified in whole-brain models, but not purely local models, cannot be explained solely as a benefit of hemodynamic modeling (although other contaminants such as motion could still be a factor). Lastly, analyses using the purely local models are equivalent to region-specific frequency-domain filtering. Although this equivalence does not imply that neural dynamics are insignificant, the signal-processing interpretation is simpler and could render the proposed neural modeling framework unnecessary (i.e. less parsimonious). Thus, the

local dynamics control serves to ensure that our guiding neural modeling framework provides additional value above its (partial) relationship to existing signal-processing techniques. This control was implemented in two distinct variants: either heterogeneous (region-specific) or homogeneous (region-invariant) autoregressive models fit to each subject.

The homogeneous model consists of an autoregressive model that is specific to subject, but not parcel:

$$BOLD_{t+1} = \beta BOLD_t + \nu_t \tag{7}$$

We assumed that the noise-component was independent and identically distributed between parcels and solved for β using linear regression (collapsing BOLD across parcels). The "input" estimates from this model consist of the residuals (ν_t). We fit the heterogenous model analogously to the homogeneous model, but with region-specific autoregressive terms:

$$BOLD_{t+1}^{(i)} = \beta_i BOLD_t^{(i)} + \nu_t^{(i)} \tag{8}$$

for parcel "i". We use these two cases to determine whether regional heterogeneity is a significant factor in any improvements due to local modeling. We refer to the homogeneous and heterogeneous models as global ("glob") and local ("loc") autoregressive (AR) models, respectively.

4. Validation and Comparison Criteria

In order to assess potential advantages of MINDy-based Filtering, we considered two types of comparisons: benchmarking (is method "a" better than "b"?), and sensitivity/robustness (how does factor "x" influence method "a" vs. "b"?). The first case establishes whether MINDy-based Filtering offers additional statistical power in detecting task effects. The second case establishes whether MINDy-Based Filtering enhances statistical power for detecting task effects in a selective (i.e., to the regions showing significant task effects to begin with) or more global manner.

4.1. Benchmarking Event-Related Effects

Trial-types were defined by high cognitive control demand vs. low cognitive control demand across the four tasks (see Sec. 3.4). Trial-specific activity was modeled using a Finite Impulse Response (FIR) model with 1TR resolution (1.2s) and task-specific length (see Sec. 3.8). Group-level statistics were compared for the peak effect (parcel $\times$ method specific) over a task-specific 2TR interval. This interval was chosen during study piloting using the peak times in conventional analyses. Thus, the analysis targets are statistically biased *against* the proposed technique since they were chosen to maximize conventional analyses. These times qualitatively

correspond with a typical HRF time-to-peak after the probe-events which define high vs. low control trials (see Sec. 3.4). Previous literature and present results suggest that these effects are primarily one-sided, with activity increased in the high-conflict (control demand) trials relative to low-conflict (low control demand) in relevant brain regions. Conversely, task-negative effects (significant decreases) have largely been associated with sustained signals as opposed to high vs. low control events. For these reasons, we only considered significant increases in activity for trial-type analyses. Group-level t-tests (within parcel) were compared for all parcels with significant increases (either method), or for a set of 34 parcels (pre-defined from independent conventional analyses which showed consistent control-demand effects across all tasks, SI Table 1, [24]). Since these parcels were pre-selected based upon conventional analyses, they are statistically biased *against* the proposed method (i.e. in favor of conventional methods).

4.2. Benchmarking Sustained Effects

In addition to event-related analyses, we also considered the identification of sustained effects (block-related changes). Results of these analyses are primarily presented in the SI (Sec. 7.3). As with event-related analyses. Sustained effects in a mixed block/event design refer to "background" activity that is present regardless of whether participants are performing a task ([37], [38]). Since we used FIR models to span each trial type, sustained effects in our analysis *only* refer to activity during inter-trial periods (non-trial periods of task-blocks) since effects during other periods are absorbed in the trial vs. rest-block contrasts ([37], [38]). We compared the group-level effect size of each technique (MINDy-based Filtering and several controls) in detecting sustained effects. Methods were compared pairwise, and benchmarking analyses were only conducted on parcels which had a significant effect for either method in a pair. Sustained analyses considered both signal increases and decreases, so the target metric was absolute t-value (1-sample group test) for the GLM sustained betas (see Sec. 3.8).

4.3. Testing Selective vs. Global Improvements

We further analyzed benchmarking results by testing how MINDy-based Filtering changes the distribution across parcels. The primary question was whether the MINDy-based Filtering: a) uniformly changes statistical power across the brain (by shift or scale); b) primarily identifies previously insignificant regions or c) primarily alters previously significant regions. This analysis is important for determining whether the technique globally improves statistical power or differentiates task-relevant regions from the rest of the brain. We test these effects using multilevel linear models to compare MINDy-based Filtering to the different control models. These

multilevel models (presented in more detail later) contain task-specific main effects of method (anatomically global) and additional terms for task-implicated (statistically significant parcels). We use these models to test the significance of model improvements (increased effect sizes) after discounting anatomically global changes.

4.4. MINDy + Hemodynamic Normalization

In addition to a naïve, linear method for estimating $h * I_t$, we considered two other variants which we present primarily in the SI (Sec. 7.5). These variants consist of implementing MINDy-based Filtering with either hemodynamic normalization or the Kalman filter. With hemodynamic normalization, estimates of latent brain effects (I_t) are calculated as usual using individualized brain/HRF models. This approach differs however in that rather than reconvolving I_t with the subject-specific HRFs, the conversion from brain to BOLD-level inferences is performed using the group-average (parcel-specific) HRF's identified from resting-state MINDy models. Thus, we leverage individualized HRF models to identify latent neural activity but reconvolve estimates with a common HRF. This normalization could be advantageous by reducing the inter-subject variation due to neurovasculature thereby better isolating common neural substrates.

4.5. MINDy + Kalman Filtering

We also tested whether leveraging the Kalman filter improves estimates of the latent input variable I_t. Kalman filtering comprises a Bayesian method to improve estimates of latent states (neural activity x_t) by incorporating priors based upon an underlying dynamic model (i.e. MINDy) and models of noise structure. Noise is divided into measurement noise (e.g. artifact) and process noise which reflects natural stochasticity in the underlying dynamical process. We considered the same stochastic-process form as before, but with I_t replaced by a stochastic variable ϵ_t:

$$x_{t+1} = W\psi_\alpha(x_t) + (1 - D)x_t + \epsilon_t \tag{9}$$

$$BOLD_t^{(i)} = [h_i * (x_\tau^{(i)} + \eta_\tau^{(i)}]_t + \nu_t \tag{10}$$

For simplicity, we again used Wiener deconvolution to produce an estimate of latent neural activity (x_t) partially obscured by noise at the neurovascular coupling (η_t):

$$y_t = x_t + \eta_t \approx h\hat{*}^{-1}BOLD \tag{11}$$

We then used the Kalman filter to estimate the latent neural activity from the deconvolved estimates. In this context, the deconvolved timeseries (y_t) functions as "measurements" in the Kalman filter terminology and the "measurement error" (η_t) corresponds to neurovascular stochasticity (i.e. errors in the neurovasculature's "measurement" of brain activity). Errors

produced by fMRI artifact (ν_t) are mitigated in the previous step of Wiener deconvolution and do not factor into our Kalman filter. We implemented the Kalman filter using statistical linearization to calculate the mean and variance of the nonlinear function ψ ([39]) using numerical sampling (n=50) at each iteration due to the lack of a simple closed-form solution. We defined the I_t estimate as the difference between the prior prediction of x_t and the posterior estimate given by the Kalman Filter (i.e. prediction errors multiplied by the Kalman gain). We estimated noise variances of 0.6 and 0.1 for ϵ and η respectively based upon resting-state data (see SI Sec. 7.1).

4.6. Sensitivity to Cognitive States

Sensitivity analyses were performed to assess the impacts of cognitive states, individual differences, and motion. In the current case, cognitive states differ between tasks and trials. Although, each of the four tasks are commonly used to index cognitive control, cognitive tasks are not construct-pure. For instance, tasks featuring delays (AX-CPT, Cued Task Switching, and Sternberg) are thought to be more dependent upon working memory than those without delays (i.e. the Stroop task). However, many task-specific factors are the same between high and low control trials of the same task (i.e. all events prior to the probe). Thus, we control for cognitive similarity across tasks by comparing results across increasing levels of cognitive similarity: low-control trials, high-control trials, and the contrast high vs. low control trials. These levels increasingly isolate the cognitive control construct by increasing control demand (high-control trials) and controlling for other task events (high vs. low contrast). Methods which are sensitive to cognitive states will produce more similar results between task contexts when the cognitive states measured are more similar. Thus, we tested whether MINDy-based Filtering increased similarity between tasks for the high vs. low contrast relative to low-control trials. We measured this using Intraclass Correlation (ICC; [40]). Tasks differed in effect magnitude and there was no theoretical basis for assuming this factor should be identical between tasks (i.e. we don't assume each task equally taxes cognitive control), so we normalized the group-average data (divided by root-sum-of-squares) for each task $\times$ method before applying ICC/generalizability analyses.

4.7. Significance Testing for Construct Identification

We used permutation statistics to compare the significance of generalizability tests between methods. When testing the generalizability of group-level patterns, we treated brain regions as the object of measurement in intraclass correlations (ICC, [40]) over task classes and estimated confidence intervals with bootstrap sampling over the set of brain parcels. We defined individual differences in terms of z-scored data relative the

group (computed task × parcel × method). We tested significance via permutation testing z-scored data that was randomly permuted between methods and then z-scored again before ICC computation (thereby maintaining the within-task distribution). Confidence intervals were computed with bootstrapped sampling over subjects.

4.8. Sensitivity to Individual Differences

We also analyzed the degree to which each method was sensitive to individual differences. This analysis also used ICC across the different levels/contrasts of control demand. However, the data of interest consisted of individual differences, which were defined by z-scores relative the group (i.e. normalized deviations from the group tendency). We computed ICCs separately for each DMCC34 parcel and determined significance using parametric statistics on the distribution over parcels.

4.9. Robustness to Motion

In an SI analysis (Sec. 7.6), we compared methods in their robustness to motion confound. While previous work has established that the model-fitting technique (MINDy) is robust to motion ([1]) it remains unknown whether MINDy-based Filtering technique also exhibits similar motion robustness. Therefore, we compared methods in terms of sensitivity to motion artifact. We considered three motion metrics for task data including the number of frames censored based upon framewise-displacement (FD) criteria ($< 0.9mm$), the median framewise displacement, median-absolute-deviation (MAD) of DVARS ([41]). We analyzed sensitivity by comparing the similarity (ICC) of results between high-motion and low-motion groups of subjects (median split for each motion measure).

5. Results

5.1. Structure and Presentation of Results

We designed analyses to answer four questions: 1) do resting-state MINDy models (partially) generalize to task? 2) does the proposed technique improve power in answering cognitive-neuroscience questions? 3) can these methods test hypotheses which were previously impractical? and 4) do improvements reflect theoretically interesting concepts (e.g. signal propagation) or do they stem from signal-processing/filtering side-effects? The first question resolves whether the intrinsic dynamics modeled at rest meaningfully generalizes to task (although not perfectly, as we are interested in the task versus rest differences). The second and third questions identify methodological contributions, whereas the last question addresses whether these techniques also offer additional theoretical insight (i.e. their success reflects some principle of brain function). This question is important for determining whether results reflect the activity-flow framework or can be

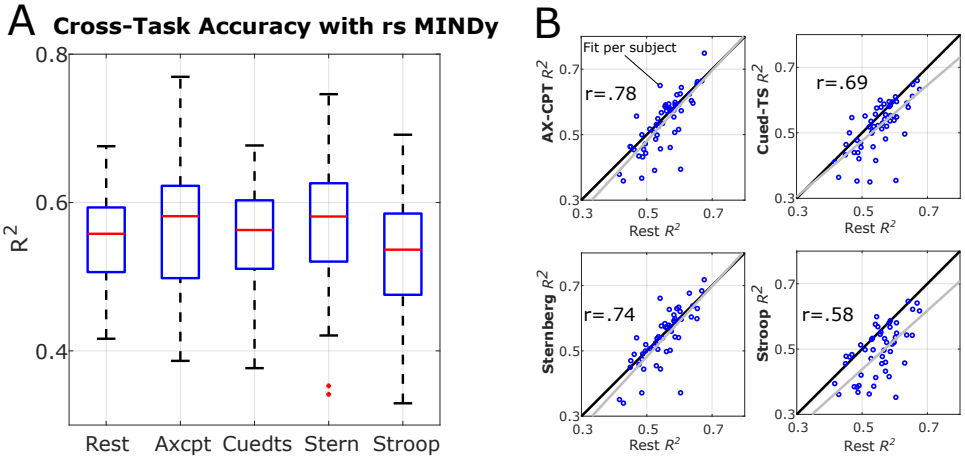

Fig 2. Resting-state MINDy models generalize to task. A) MINDy models trained on resting-state data produce similarly accurate predictions for task data. Goodness-of-fit is quantified in terms of the mean R^2 value across all brain parcels ([31]) and scanning runs (n=6) in predicting the difference time series: $X_{t+1} - X_t$ ([1]). B) Individual-differences in model accuracy are highly correlated between resting-state and task data.

more parsimoniously explained in terms of (non-neural) signal processing effects.

In the main text, we emphasize comparing methods in event-related analyses due to the popularity of event-related designs. However, we also compared methods for the analysis of sustained-effects in a mixed block/event design. These results are presented in SI Sec. 7.3 and 7.4. We also tested the specific contribution of modeling connectivity by comparing MINDy-based Filtering with analogous filters using reduced (autoregressive) models (SI Sec.7.2).

5.2. Resting-state Model Predictions Generalize to Task

We first test whether the proposed technique actually serves as a conceptual "filter" in removing intrinsic-dynamics from task, rather than making this information more salient (as would occur when there is little overlap between task and rest). Our framework assumes systematic discrepancies between task and resting-state (i.e., we are interested in the difference between contexts), but we assume that there is some overlap of task and resting-state dynamics for us to "remove". In statistical terms, we first ensure that the approach removes variation from task data (associated with covariance of task-rest dynamics) rather than adding additional variation as would occur when subtracting independent factors. To test this possibility, we first quantified the goodness of fit for resting-state models (MINDy; [1], [17]) in predicting both task and resting-state data within subject. Paired (within-subject) analyses indicated greater goodness-of-fit for resting-state

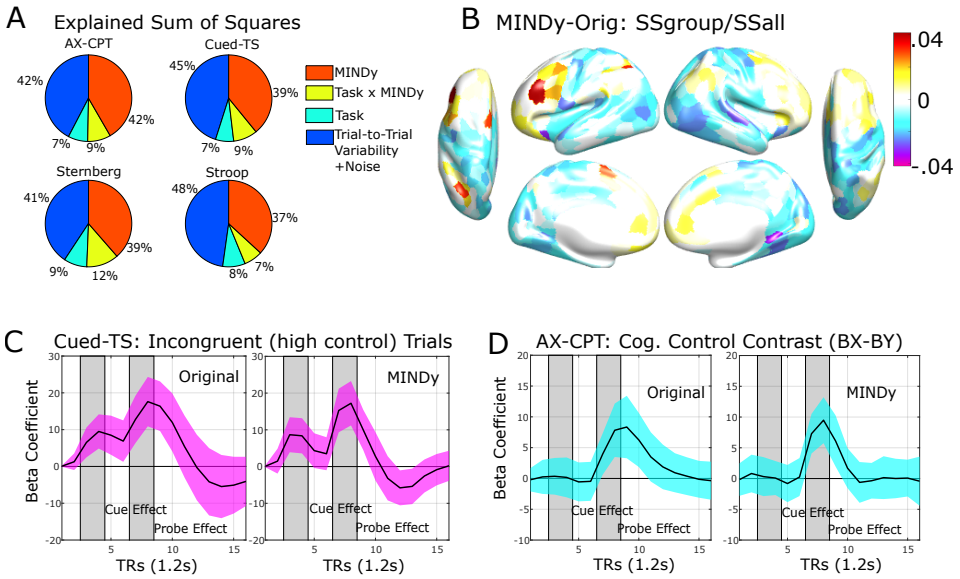

Fig 3. MINDy-based Filtering reduces variability within and between subjects. A) MINDy-based Filtering accounts for a significant portion of unique variability within each subject's data. This effect holds across tasks (results averaged over all parcels, subjects). Variance partitioning was performed after removing variation due to nuisance factors (motion and drift). B) Difference in the relative group-explained variability between MINDy and the original data. Note that MINDy-based filtering actually decreases the proportion of group variance in some regions, but increases for task-implicated regions (e.g. lPFC). C) MINDy-based Filtering reduces the between-subject variability of task-evoked signals. Example shown is the mean signal over the DMCC34 parcels for the Cued-TS high control-demand condition (incongruent trials). D) Variability also decreases for contrasts between conditions. Example shown is for the AX-CPT (BX-BY contrast).

scans ($R^2 = .55 \pm .06$) relative to task scans (mean R^2 across tasks=$.52 \pm .08$, paired-$t(52) = 4.37, p = .5.9E-5$; Fig.2A), but the magnitude of this difference was not particularly large ($\Delta R^2 = .03 \pm .05$). Moreover, individual differences in goodness-of-fit were consistent across tasks (Fig. 2B), which indicates that model accuracy is also highly preserved within-subject. We conclude that the short-term evolution of brain activity is similar (but not identical) in resting-state and various task contexts. By leveraging large-scale resting-state models (MINDy) the proposed technique filters out intrinsic dynamics common to resting and task state.

5.3. MINDy-based Filtering Accounts for intra and inter-subject Variability

We also tested whether these intrinsic dynamics explain unique variability above the task GLM. This test is important for determining whether MINDy serves to predict the mean brain-response for each trial-type or whether it also predicts trial-to-trial variability. We quantified these properties through sum-of-squares partitioning (ANOVA). Across all tasks, we found that the proportion of unique variance explained by MINDy was significant

(39.1% on average, Fig. 3A). However, MINDy predictions and the task effects do have some overlap (a non-zero MINDy×task sum-of-squares, Fig. 3A), thus MINDy predictions account for some of the variation in both the trial-to-trial variability (variation unique to MINDy) and the typical response across trials (MINDy × task interaction). We also tested how MINDy-based Filtering impacts variability in the evoked-response between subjects. We restricted these analyses to the pre-defined set of regions (the DMCC34 parcels, [24]) which were previously identified as having a significant control-demand effect across tasks. Results demonstrated that MINDy filtering decreased inter-subject variability in both main effects of trial-type (e.g. Fig. 3C) and the contrast between trial-types (e.g. Fig. 3D). In particular, these analyses and associated event-related timecourse visualizations reveal that the peak task-related effects become sharper (more well-defined), as well as more temporally-precise, after MINDy-based filtering. We used ANOVA to partition variance in the cognitive control effect into group-level variance and individual variance over the relevant trial periods.

We then tested whether MINDy increased the proportion of cognitive control effects attributed to a common group factor (sum-of-squares for the group effect divided by the total effect). As expected, regions implicated in cognitive control, such as the lateral and medial prefrontal cortex, anterior insulae, and posterior parietal cortex, had larger proportions of variability explained by the common group factor (analogous to Fig. 4A). Interestingly, although MINDy-based Filtering increased group variability (decreased inter-subject variability) in task-implicated regions, it decreased the common group factor for regions not implicated in cognitive control (Fig. 3B). The proportion variance explained by a common group effect increased across the DMCC34 parcels (paired-$t(33) = 3.91, p = 4.3E-4$). Thus, by removing individual-differences in intrinsic brain dynamics, MINDy-based Filtering reveals more similar task-effects between subjects.

5.4. Improved Group-Level Detection Power

We tested whether MINDy-based Filtering improved statistical power in detecting group-level neural effects for each task, and in an omnibus test across tasks (Fig. 4A,B). For each event-related pairwise comparison of methods, we tested the change in effect-size (group t-value) for parcels demonstrating a significant increase ($p < .01$, uncorrected) for either method within a pair. Results indicate that MINDy-based Filtering significantly increased statistical detection power on all tasks (four of four) for the event-related contrast relative to controls (all p's≤1.5E-7; Fig. 4 B). For omnibus analyses, we collapsed observations across tasks (Fig. 4A). Results indicated that MINDy-based Filtering generally increases detection

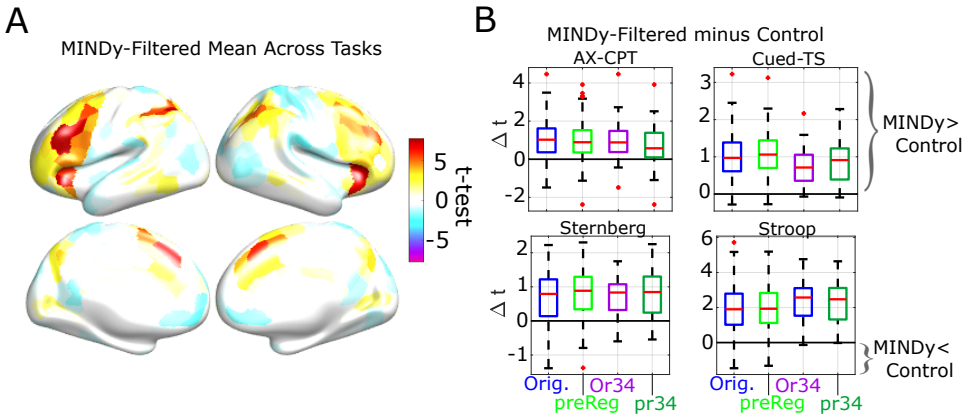

Fig 4. MINDy-based Filtering increases power in detecting event-related task effects. A) Spatial distribution of high-vs.-low conflict effects averaged across tasks using the proposed technique. B) Paired comparisons of effect size between control methods and the proposed approach in identifying high-vs.-low conflict effects for significant brain parcels ($p < .01$). Values greater than zero indicate that the propose technique improves upon controls. Comparison techniques are denoted:"Orig"=original analysis (no filter), "prReg"= motion and drift pre-regressed before GLM fitting. "34" references analyses constricted to the DMCC34 set of parcels (see SI Table 1, [24]), rather than all significant parcels. Or34 compares these parcels after MINDy filtering to the original analysis and pr34 to analyses with the pre-regressed data.

power for event-related analyses relative to controls (vs. original: paired-$t(541) = 28.3$, $p \approx 0$, vs. pre-regressed: $t(540) = 27.8$, $p \approx 0$). We conclude that the proposed techniques improve group-level detection of task effects.

One limitation of the previous tests, however, concerns the determination of which parcels are included in analysis: we compared effect sizes in parcels that were statistically significant (i.e., large effect sizes). This approach is anatomically parsimonious in that the comparison regions are informed by data rather than prior assumptions. However, this dependency could produce biases due to differences in higher-order features (e.g. overdispersion) between methods. Therefore, we repeated the previous analyses over a fixed set of 34 pre-specified brain parcels (SI Table 1, [24]) that demonstrated significant increases due to cognitive conflict (event-related contrast) across all four tasks during independent and pre-specified analyses (see Methods, [24]). The implicated parcels agree with previous studies mapping the neuroanatomy of cognitive control and are largely located along lateral prefrontal cortex and anterior insula (Salience/Ventral Attention and Control networks; [42], [31]). Analyses over this restricted, pre-specified group of parcels agreed with the previous results: the omnibus (all task) statistical detection power improved relative all controls (maximum $p = 4.6E-4$; Fig. 4B). Thus, results indicated that MINDy-based Filtering improved statistical detection even when analyses were restricted to this group of 34 pre-specified parcels.

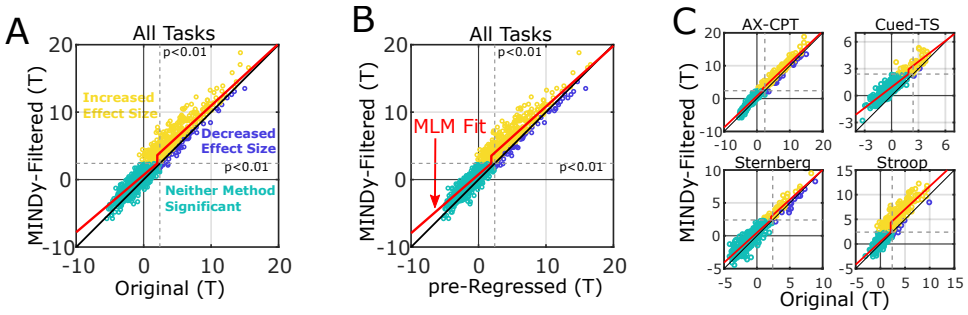

Fig 5. MINDy-based Filtering enhances task-related signals relative to controls. A) Comparison of parcel significance before and after MINDy-based Filtering collapsed across tasks. The multi-level model fit (averaged across the main effect of task) is plotted in red and the threshold-nonlinearity indicates sensitivity to parcel-significance. B) Same as (A) but for the pre-regressed control. C) Task-specific comparisons relative the original analyses. Improvements can be seen in the number of parcels exhibiting higher t-values after MINDy-based Filtering relative to conventional analyses (i.e., above the identity line). Yellow dots indicate significant parcels (in terms of the control-demand effect) which also had increased effect sizes from MINDy-based Filtering, while blue dots denote significant parcels whose effect sizes were larger with conventional analyses. Teal dots denote parcels which did not exhibit a significant control-demand effect for either method.

5.5. MINDy-based Filtering Selectively Enhances Task-Related Neural Signals

Results in the previous section indicate that the proposed technique increases the statistical detection power of task effects (Fig. 4B). Statistical power and effect sizes are useful benchmarking criteria as they are easy to interpret and relate to potential applications. However, these markers are also limited in that they indicate the ability to reject a generic null hypothesis of no task effects. Yet this generic null is not always a useful benchmark from which to provide additional scientific insight. For instance, approaches which magnify anatomically global effects may provide little benefit to functional "brain-mapping" studies, which are most meaningful when they differentiate between brain regions. Therefore, we tested whether the improvements found with MINDy-based Filtering are anatomically global or serve to further differentiate regions (i.e., are anatomically selective).

We consider two sorts of global effects: additive "shifts" in the global signal and global "scaling" of task effects. In statistical modeling terminology, the former reflects a main-effect (intercept) of method, whereas the latter reflects the method-specific slope. We modeled the differentiation between brain regions as either a main effect of regional significance (i.e., whether a region has a significant effect) or as an interaction with regional significance reflecting either a shift or rescaling of effect sizes of significant regions due to MINDy-based filtering, relative to the control models. We use the logical-valued variable $Sig_{task,Parc}$ to denote whether a parcel exhibits a

significant effect for either method in a given second-level task analysis. We denote the MINDy-filtered second-level estimate (group-T) for each as $X_{task,Parc}$ which is modeled as a function of matched control analyses (e.g. the original GLM or pre-Regressed) which are denoted $Y_{task,Parc}$:

$$X_{task,Parc} = \beta_{task} + \beta_0 Y_{task,Parc} + Sig_{task,Parc}(\beta_1 + \beta_2 Y_{task,Parc}) + \epsilon_{task,Parc}. \tag{12}$$

We assume that ϵ is independently and identically distributed across tasks and parcels (iid.). The coefficient β_1 represents the main effect of parcel significance as a binary factor, while β_2 represents the interaction with parcel effect size in control methods. Conceptually, these two components represent the degree to which MINDy-based Filtering further separates task-implicated and non-implicated parcels and the degree to which differences among task-implicated regions are further magnified, respectively.

Results indicate that the MINDy-based Filtering technique demonstrates differential sensitivity, in that improvements are greater in task-implicated regions (Fig. 5A, B). The main effect of event-related regional significance was significant for all controls (vs. original and vs. pre-Reg: $t(1669)$=13.34, $p \approx$0, t=14.36, $p \approx$0, respectively). This result indicates that MINDy-based Filtering further separates event-implicated and non-implicated regions rather than simply increasing global statistical features. This feature also held at the single-task level in which linear models revealed a main effect of regional significance in all four tasks for both original and pre-regressed controls (min. $p = .007$; Fig. 5C). MINDy-based Filtering also differentially magnified effect sizes relative the original analysis ($t(1669) = 3.3, p = .001$), but this effect was small and did not reach significance for the pre-regressed control ($t = 1.51, p = .13$, 2-tailed). Thus, task-implicated regions experienced the greatest improvements due to MINDy-based Filtering. For the current dataset, this approach primarily functioned to further highlight task-implicated brain regions (a main effect of regional significance) rather than magnifying the differences between task-implicated regions. These results imply that MINDy-based Filtering is sensitive to task-implicated brain regions rather than inducing anatomically global effects.

5.6. Advanced estimates for I_t could improve performance

In the present work we have emphasized the power of a simple differencing approach to estimate I_t, namely by taking the difference of observed BOLD and predictions based upon the resting-state MINDy model. However, we also considered two variations on this approach: one in which estimates for I_t are normalized based upon the group-average HRF for each parcel (estimated with resting-state MINDy) and another in which I_t was estimated using the Kalman Filter (which uses Bayesian estimation). We found cases

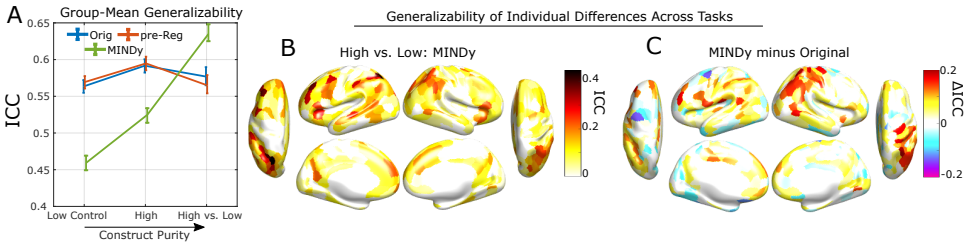

Fig 6. MINDy-based Filtering improves identification of cognitive control demand. A) Cross-task generalizability (ICC) for group means in the low-conflict condition (low construct purity across tasks) and the high vs. low conflict contrast (high purity) by method. B) Anatomical distribution of generalizability for the proposed technique and C) The difference in cross-task generalizability between MINDy-based Filtering and conventional analyses. Comparison techniques are denoted:"Orig"=original analysis (no filter), "Pre-Reg"=motion and drift pre-regressed before GLM fitting, and "MINDy"=MINDy-based Filtering.

in which these approaches further improved the detection of group-level effects above the benefits conferred by naïve MINDy-based filtering. A full account of these alternatives is detailed in the SI (Sec. 7.5). However, we chose to emphasize the naïve MINDy-based filter for the main text so that improvements could be more directly linked to the underlying neural framework (i.e., input to a dynamic system) since this filter does not add any additional mitigation of nuisance factors. By contrast, improvements in the other estimators might be due, in part, to the mitigation of nuisance factors such as inter-subject variation in neurovasculature (HRF normalization) and "noise" factors (Kalman filter).

5.7. Identifying Individual Differences in a Latent Cognitive Construct

The previous analyses indicate that MINDy-based Filtering enhances the identification of neural activity associated with a set of contrasts between trial-types (theoretical high control-demand trials minus low control-demand trials). However, many cognitive neuroscience studies seek to understand cognitive constructs, as opposed to unitary tasks. In the current section, we explore how well each method identifies the neural correlates of one such construct: cognitive control. The four tasks we studied have all been previously used to index cognitive control (typically via the difference between high-control and low-control trials). However, because the tasks themselves are not construct-pure (i.e., they tap multiple cognitive constructs) the neural activity associated with tasks is also expected to be non-identical. To control for this fact, we used the different trial types to generate levels of "construct-purity" in terms of cognitive control: low-control trials (low purity) and the high-vs.-low contrast (high purity). We consider the high-vs.-low contrast to be more "construct-pure" in terms of cognitive control since it controls for many of the other cognitive processes that differentiate tasks. For instance, speech production (unique to the Stroop

task), is identical between high and low-conflict trials (the same set of words are produced). Likewise, working memory maintenance during delays (Sternberg, AX-CPT, and Cued-Task Switching) does not differ between high and low control-demand trials since these trial-types are identical through the delay period (up until the probe). The high-control trials form an intermediary between the "high-purity" high-vs.-low contrast and the "low-purity" low-control trials. We use this condition for illustration (i.e. it suggests a continuous relationship in Fig. 6A) but did not perform statistical testing as we did not have any explicit hypotheses regarding this condition.

We tested how sensitive each approach was to the cognitive control construct via the relationship between "construct-purity" and cross-task similarity of neural effects. For this test, we indicate that a measure is "sensitive" to a factor (cognitive constructs) if the similarity in measurements reflects the similarity in that factor. We therefore consider a measure "sensitive" to cognitive constructs if it reports higher similarity between tasks for the high "construct-purity" condition (high-vs.-low control demand contrast) than for the low "construct-purity" condition (low demand trials).

We quantified "similarity" across the four tasks using Intra Class Correlation (ICC, [40]), and performed analyses in terms of both the group-average and individual-differences. For group-average analyses, ICC "units of observation" consisted of the mean beta for each brain parcel (all 419 brain regions) and "classes" consisted of the different tasks. Results indicated that the proposed technique was sensitive to the cognitive control construct at group level (Fig. 6A). In the "low-purity" condition, with MINDy-based Filtering there was significantly lower similarity between tasks ($ICC = .46 \pm .01$) than conventional approaches (both $p's < .001$, 5,000 bootstraps). Thus, MINDy-based Filtering does not generically increase the similarity of task results irrespective of cognitive construct. By contrast, for the "high-purity" condition, MINDy-based Filtering generated significantly more similar results across tasks ($ICC = .63 \pm .01$) than conventional analyses (all $p's < .001$, 5000 paired bootstraps). We conclude that MINDy-based Filtering improves sensitivity to the cognitive control construct at group-level. Based on the nature of how these ICCs were calculated, the finding can also be interpreted as indicating that the anatomical profile of effects (i.e., the gradient of effect sizes across the brain) becomes more similar or consistent across tasks after MINDy-based filtering, relative to conventional analyses, and the controls.

We next tested the generalizability of individual differences. For these analyses we normalized data between subjects for each method $\times$ task $\times$

parcel before computing ICC separately for each parcel (using subjects as the "units of observation"). In general, ICC values were highest for task-implicated regions (e.g. lateral prefrontal cortex, anterior insulae and posterior parietal cortex; Fig. 6B) and the MINDy-based Filtering demonstrated particularly high generalizability at posterior parietal cortex relative to conventional techniques (Fig. 6C). We tested the generalizability of individual differences restricted to the DMCC34 set of parcels. Results indicated that MINDy-based Filtering improves identification of individual differences over conventional techniques for the high vs. low control contrast (vs. original: $paired - t(33) = 2.47, p = .019$; vs. pre-regressed: $t = 3.56, p = .001$). As with the group-level analysis, generalizability did not increase over conventional analysis for the low construct-purity (low demand) condition. Thus, MINDy-based Filtering improves the estimation of neural individual differences related to a cognitive construct but does not artificially increase generalizability across cognitively dissimilar task conditions.

These results also offer (speculative) theoretical interpretations. The inter-task variability of conventional techniques and MINDy-based Filtering are linked since "input" is defined as the difference between the observed brain activity and the propagation along intrinsic dynamics ($x_{t+1} = f(x_t) + I_t$). Consequently, this section's results suggest that cognitive control signatures are most similar between tasks when they first impact the brain (as "inputs") but lead to more task-specific patterns of activity (x_t) as they evolve (or propagate) according to intrinsic brain dynamics. Analogous results suggest that individual differences follow the same principles. Of course, these interpretations are post-hoc and mainly serve to demonstrate the potentially utility of MINDy-based Filtering. Future studies should explore these possibilities in more detail.

5.8. MINDy-based Filtering Enhances Brain-Behavior Relationships

The previous section demonstrated that neural effects identified with MINDy-based Filtering better generalized across task conditions tapping a common construct (cognitive control). In this section we demonstrate that this relationship also holds for behavior by using individual differences in task effects to predict the corresponding variation in cognitive control effects. These effects are measured by the difference in behavioral response time (RT) and accuracy during high-control vs. low-control trials. For each subject$\times$ task $\times$ session, we summarized event-related effects in each task$\times$ method via the difference of normalized (z-scored over subjects) high and low control trial coefficients averaged over the DMCC34 set of parcels and similarly for sustained effects. Behavioral measures were similarly defined by the difference in normalized RT and accuracy between high and low control trials. Interestingly, we found that across methods, individual

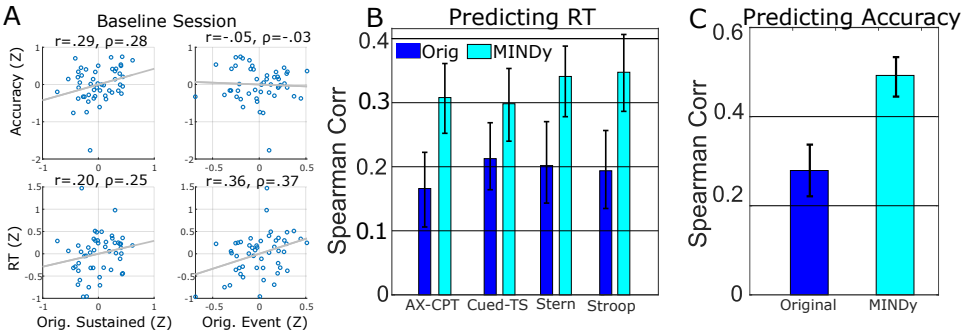

Fig 7. MINDy-based Filtering enhances prediction of individual differences in cognitive control effects. Behavioral measures correspond to the difference of high control trials minus low control trials in normalized response times (RT) and accuracy (Acc). A) Individual differences in DMCC34 sustained responses predict accuracy while event-related responses (high vs. low control contrast) predict RTs in the baseline session (collapsed across tasks). Data plotted correspond to the original (non-MINDy) analyses. B) MINDy-based Filtering enhances correlations between event-related responses and RTs in each task (collapsed across the 3 scanning sessions). C) MINDy-based Filtering also enhances the correlation between sustained responses in DMCC34 and error-rates (baseline session). Confidence intervals correspond to 68.3% (i.e., $\pm$ 1 SD in a normal distribution)

differences in RT were positively correlated with the conflict-related (event) brain response but had a weaker relationship to sustained activity (Fig. 7A). By contrast, individual differences in accuracy were positively correlated with sustained activity, but unrelated to event-related activity (Fig 7). Therefore, we compared methods in predicting RT using event-related estimates and in predicting accuracy using estimates of sustained activity. We found that MINDy-based Filtering increased the recorded correlations with RT for each task (Fig. 7B) and the change in mean correlation across tasks was statistically significant (vs. original, pre-regressed: $p < .05$, 5,000 bootstraps). Similarly, our approach increased correlations with accuracy ($p < .05$, 5,000 bootstraps, Fig. 7C). Results using the pre-regressed control are depicted in SI Fig. 10. We conclude that after MINDy-based Filtering, individual differences in brain responses better predict behavioral measures associated with cognitive control.

6. Discussion

We demonstrated that MINDy-based Filtering increases the ability to detect both event-related (cognitive control-demand) and sustained brain responses in task fMRI (Sec. 5.4, SI Sec. 7.3). These effects are strongest in task-implicated brain regions (Sec. 5.5) and generate higher temporal precision than the original BOLD timeseries. By accounting for intrinsic dynamics, MINDy-based Filtering accounts for both trial-to-trial variability within subjects, and variability between subjects (Sec. 5.3). However, while the absolute magnitude of subject-to-subject variability decreased,

individual differences (and group—level activity) in a latent cognitive construct (control-demand) generalized better between tasks after MINDy-based Filtering (Sec. 5.7). MINDy-estimated task effects were also more predictive of individual differences in behavior (Sec. 5.8). These results suggest that MINDy-based Filtering can enhance the detection of task-evoked brain activity.

6.1. Relationship with Frequency-Based Filtering

Frequency-based (spectral) filtering has been applied to fMRI signals in many previous studies ([43],[44]). High-pass filtering is commonly applied to both resting-state and task data to remove signal drift which is thought to largely reflect changes in non-neuronal variables. Low-pass filtering is also sometimes applied, primarily for resting-state data. Although these approaches were common in early fMRI experiments, the changing nature of fMRI acquisitions (e.g. TR length) and analyses (e.g. functional connectivity) has led to renewed debate over these techniques ([45]) and the development of more sophisticated methodologies (e.g. [46],[47]). In the current work, we did not perform spectral filtering (instead using AFNI's "polort" function for polynomial basis de-drifting). Likewise, MINDy-based Filtering is not a direct replacement for spectral filtering, which can be applied before our technique, afterwards or not at all. However, as previously mentioned, when the connectivity parameter of our model is zero, the proposed technique reduces to a form of spectral filtering based purely upon autoregressive models. Empirically we have demonstrated that MINDy-based filtering outperforms filters based upon autoregressive models (SI Sec. 7.2, SI Fig. 8 A), so effects cannot be attributed solely to removal of particular frequency components within each region.

Notably, MINDy-based Filtering improves detection in both sustained and event-related analyses over both conventional methods and autoregressive filters. By contrast, filters based upon autoregressive models are expected to underperform in the identification of (low-frequency) sustained effects as we confirmed in our analyses (Sec. 5.4, Fig. 4D). At a statistical-level, dynamical systems models (including MINDy) capture the multivariate partial autocovariance between successive time-points (i.e. how x_{t+1} is related to x_t). As a result, removing these predictions from the training data (Rest) inherently yields a timeseries with lower autocovariance. The improved detection of sustained effects is therefore significant as it indicates that the proposed technique reveals systematic differences between the resting-state and task dynamics rather than simply acting as a high-pass filter. These effects are also more pronounced in task-implicated parcels (Sec. 5.5, Fig. 5A-C) indicating that these features are also context-related.

6.2. Relationship with other approaches

The current approach is conceptually related to several current initiatives for linking resting-state and task-state brain activity. Our approach uses resting-state brain dynamics to extrapolate patterns of intrinsic dynamics that also factor into brain activity during task states. Frameworks such as Activity Flow ([8]) have demonstrated similarity between the spatial aspects of evoked responses and resting-state network structure. Likewise, functional connectivity patterns have been found to be roughly similar between resting-state and task ([48]). However, whereas these frameworks are largely employed to discover similarities between spontaneous and evoked activity, we analyze the manner in which task-state deviates from resting-state activity over short time-scales.

Other approaches have also investigated the difference between brain dynamics in task-state and resting-state. Previous work ([11], [10]) has demonstrated that intrinsic dynamics shape task-evoked activity on a trial-by-trial basis and modeling studies have reproduced the statistical differences between task and resting-state activity ([12]). Our approach furthers these efforts by leveraging these underlying concepts into an empirical modeling/analysis framework.

Dynamic Causal Modeling (DCM, [19]) frameworks have also used empirical dynamical systems models to improve estimates of task effects. As previously mentioned (Sec. 1.2), DCM techniques allow task effects to manifest changes in the exogeneous drive to brain regions and (for small-scale DCMs) the effective coupling between brain regions. By contrast, the current MINDy-based Filtering technique only models a single factor: changes in the input to each brain region, which collapses both of these mechanisms into a single term as is common in larger-scale DCM models (e.g. [21]). Our approach differs from all DCMs, however, in two fundamental ways.

First, whereas DCMs fit all data simultaneously, we parameterize our dynamic (MINDy) models solely from resting-state data. As a result, our model parameters are not impacted by any preconceived models of task effects (i.e., that they follow a certain temporal pattern). Secondly, we do not explicitly model task effects. Whereas DCMs directly fit parameters to task conditions, MINDy-based Filtering produces a full timeseries of estimated effects based solely upon the observed fMRI scan data (i.e. no task information is used). Thus, our approach estimates the evolution of latent variables (task-related "input" to each region) rather than estimating coefficients for a pre-specified temporal model of task effects. As such, MINDy-based Filtering is much more flexible than DCM, as it functions as a processing step rather than a full analysis pipeline in and of itself. In the

current work, we used statistical GLMs to analyze the MINDy-filtered data with Finite Impulse Response models fit for each trial type and additional components to model task blocks (mixed block-event design). However, the end-product of our technique (a timeseries) could, in principle, be analyzed with a wide variety of methods, including parcel-level multivariate techniques (e.g., multivariate pattern analysis; MVPA). Thus, although DCM and MINDy-based Filtering both use empirical dynamical systems models with similar assumptions, the approaches differ radically in how these models are leveraged.

6.3. Limitations

The proposed work rests upon three related claims: 1) intrinsic dynamics are roughly conserved between task periods and rest, 2) that by subtracting intrinsic dynamics we identify changes in "input" to each brain area and 3) that the signal generated by this calculation is a better marker of task effects (ostensibly task-related cognition). The first two claims are interdependent. We have mathematically defined changes in "input" as the signal components which are not explained by intrinsic dynamics (the residual after subtracting the modeled intrinsic component). The accuracy of estimated changes in "input" thus hinges upon whether the modeled intrinsic dynamics meaningfully generalize. We attempted to address this question empirically (see Sec. 5.2), and the results suggest that this assumption does hold. We also note that the generalizability assumption is "soft" in the sense that small changes in effective connectivity do not violate our assumptions. Since each connection describes the strength of input to the "post-synaptic" region, changes in connection strength are absorbed in the input estimate (summing over "pre-synaptic" sources). However, we do assume that our filter removes variance which could be violated by some forms of large, systematic changes in effective connectivity. Fortunately, this assumption is easy to check (e.g., see Sec. 5.2) and we have not found evidence of its violation.

6.3.1. Methodological Considerations

The bulk of our results concern the last claim (improved detection power) and the demonstration that observed statistical improvements are related to task-specific neural processes. We performed these tests using several controlled comparisons and lines of inquiry. However, our efforts in this domain are limited by using a specific subset of cognitive tasks: those used to index cognitive control. As the set of potential cognitive constructs remains vast, further testing in other cognitive domains may be useful.

Another limitation concerns how MINDy models are parameterized. Since we parameterize models based upon resting-state data, we require the collection of both resting-state and task data for each subject which

increases data requirements. Moreover, this dependency could prove problematic for low-quality resting-state data, as mis-specified resting-state models could corrupt task estimates. We found that individual differences in goodness-of-fit were consistent across tasks (Fig. 2B) so this possibility cannot be ruled out. However, previous analyses of MINDy modeling indicated that the goodness-of-fit is not related to individual differences in motion ([1]). The results also do not support model overfitting, as goodness-of-fit did not uniformly decrease when applied to novel (task) data relative to training (rest) data (Fig. 2A). Further study may therefore be beneficial in determining which factors (neural or nuisance) influence individual differences in goodness of fit, as these factors could influence estimated individual differences in task variables.

6.3.2. Mechanistic Considerations

Future study is necessary is necessary to disambiguate which biological mechanisms contribute to the calculated "input" signal. For decades, computational neuroscience models have largely formalized task context as an exogeneous forcing ("input" or "bias") term in neural networks and connectionist models (e.g. [49], [50], [51], [52], [53]). This formulation is appealing for its simplicity; however, external contexts serve only as "inputs" during sensory transduction, since brain activity is known to modulate even sensory neurons (e.g. [54], [55]). Even when these effects are neglected, many modeling studies assume that brain regions receive task "inputs", even if these regions are not directly enervated by sensory nerves (e.g. [51]). As a result, these "inputs" should not be interpreted as literal inputs to the brain (i.e. signals from sensory nerves). Rather, these "inputs" include the initial propagation of such signals over the fMRI sampling rate (1 TR), so our approach is limited by the temporal resolution of fMRI BOLD.

The nature of these "inputs" is also somewhat underspecified. In the current approach, we use MINDy to model the propagation of brain signals during resting-state. The model predicts task-fMRI activation based upon the effective connectivity parameters estimated from resting-state. However, these parameters are limited to describing the relationship of bulk activity between brain regions. Many brain regions contain diffuse sets of neurons with heterogeneous axonal connectivity profiles. Several lines of evidence suggest that task-contexts can modulate the effective connectivity between brain regions via selective recruitment of neurons in synchronous ensembles ([56], [57], [58]). Our approach is therefore limited, in that it does not explicate how changes in "input" relate to changes in the effective coupling between brain regions. Future studies may improve upon the current approach by further modeling how task events modulate effective connectivity between brain regions. Such studies could either directly pa-

rameterize connectivity × task interactions (as in DCM), or extend the filtering approach to estimate time-varying (or state-varying) connectivity.

6.4. Task Dynamics Could Potentially Influence Statistical Improvements

The current approach serves to estimate latent changes in input to each brain area. In the present study we found that MINDy-based Filtering consistently improved statistical detection power across tasks. However, there may be contexts in which brain activity ($x(t)$) is a more consistent marker of task context than input ($I(t)$). Such cases occur when different input patterns (i.e., inter-trial variability in input) lead to the similar outcomes in terms of activity. In these cases, MINDy-based Filtering might actually decrease detection power, since the "input" on each trial is less consistent than its long-term consequences. Future studies might identify such cases using a wider variety of tasks.

One area in which our approach could also be limited is in detecting slow neural events in which task-related activity evolves over multiple TRs. Since our approach acts as a pre-processing filter (i.e. doesn't use task information) it is possible that it could filter out the propagation of very slow task-related activity in addition to task-unrelated activity. However, this cancellation is only expected when task-related activity propagates identically (has the same dynamics) to spontaneous brain activity. In practice, we have found that MINDy-based Filtering improves the detection of sustained brain activity and strengthens brain-behavior linkages (Sec. 5.8, SI Sec. 7.3).

6.5. Conclusion

In the current work, we proposed a new technique to estimate the influence of external contexts (task conditions) on brain activity (in our case fMRI). This technique forms a mathematical filter and therefore functions as a preprocessing step rather than as a direct tool for hypothesis testing. This property is advantageous as it allows this approach to be used in conjunction with a variety of existing methods. We have demonstrated that using MINDy-based Filtering improves statistical power (Fig. 4), increases sensitivity to task-implicated regions (Sec. 5.5; Fig. **??**)), and better identifies the neural signatures of a latent cognitive construct (cognitive conflict) in both individuals and group-level (Fig. 6). Moreover, MINDy-based Filtering enhances the strength of brain-behavior reslationships (Fig. 7). These improvements are not sensitive to motion within a reasonable range (SI Sec. 7.6). Our technique can be easily inserted into most fMRI processing pipelines and we have made code available via the primary author's GitHub to facilitate this process.

7. Supplemental Information

7.1. Estimating Error Variances for the Kalman Filter

We implemented a nonlinear Kalman Filter using numeric statistical linearization ([39]) of the nonlinear function ψ in computing expected values and variances. Kalman Filtering was performed in the deconvolved space (see Sec. 4.5) so the "measurement" noise consisted of stochasticity in the neurovascular coupling. To avoid biasing results, we assumed that both process (neural) and measurement (vascular) noise were identically and independently distributed between brain regions which left two remaining scalar parameters: the measurement (R) and process (Q) noise variances. To estimate these quantities, we first constrained the total contribution of all stochastic terms (estimation error, measurement noise and process noise) to roughly match the prediction error variance in the deconvolved space (.6-.8 depending upon parcel). For three random subjects we finely sampled the 2D space of Q,R variances (valued between 0 and 1) and compared Kalman filters with the associated parameters in terms of prediction accuracy (log-likelihood of measurements given predictions) for resting-state scans. These results generally favored a small value for R relative to Q but were relatively smooth. We then compared prediction accuracy over task scans for all subjects using three candidate value-pairs: $\{Q = .6, R = .1\}$, $\{Q = .3, R = .5\}$, and $\{Q = .5, R = .5\}$. As expected, we found the best predictions (highest likelihood) for $\{Q = .6, R = .1\}$ on all tasks, so we used these values for further analyses. However, we note that the other two candidates produced more accurate estimates of individual response-times on the Cued-TS and Sternberg tasks than the chosen values (see Sec. 5.8), whereas all other properties were nearly identical.

7.2. Comparison with Reduced Models

We compared estimation of inputs using MINDy models to analogous estimates to reduced autoregressive forms with autoregressive terms which were either subject-specific (but not parcel-specific) or terms which were specific to subject and parcel (see Methods Sec. 3.10). Since the MINDy model also features an autoregressive term (the "Decay"), these alternative models serve as reduced special cases which don't include the effects of inter-regional signaling (connectivity). As such, improvements of the full MINDy model over these alternative (autoregressive) models indicate the contribution of modeling connectivity, as opposed to simply accounting for purely local dynamics.

Results indicated that group-level detection power for MINDy-based Filtering was greater than both the homogeneous and heterogeneous autoregressive comparison models. Detection power was greater across all four tasks for both event-related (maximum $p = .028$, Fig. 8A,B) and

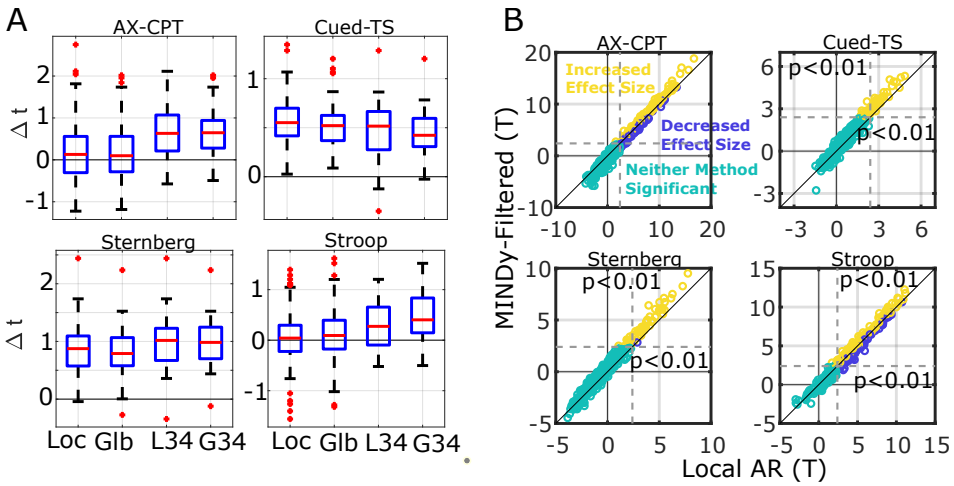

Fig 8. MINDy-based Filtering provides greater detection power using the full model over autoregressive (AR) reduced models which do not model connectivity. Comparisons are for the high-control vs. low-control contrasts. A) Pair-wise difference in detection power (group T) for significant parcels and the DMCC34 parcels by task and model. "Loc" denotes the locally-parameterized (heterogeneous) AR while "Glob" denotes the globally-parameterized (anatomically homogeneous) AR model. The suffix "34" denotes when the DMCC34 parcels were used as opposed to all significant parcels. B) Scatterplots of parcel significance when filtering with the local AR model vs. full MINDy model for each task. Yellow dots indicate significant parcels (in terms of the control-demand effect) which also had increased effect sizes from MINDy-based Filtering, while blue dots denote significant parcels whose effect sizes were larger with conventional analyses. Teal dots denote parcels which did not exhibit a significant control-demand effect for either method. Note that for all tasks, the most significant parcels (top right) benefit from using the full MINDy model even when the average effect over all significant parcels is small (e.g. AX-CPT and Stroop in A).

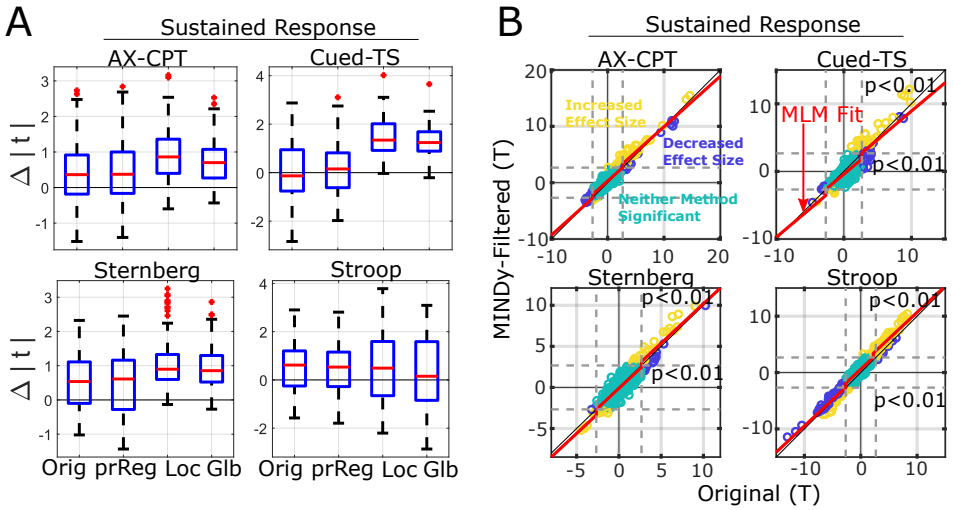

Fig 9. MINDy-based Filtering generally improves the detection of sustained effects. Unlike event-related effects, we permitted bidirectional sustained effects hence we compared the absolute magnitude of group-T statistics. The definition of significance was likewise 2-tailed. A) Pair-wise difference in detection power (group T) for significant parcels by task and control ("Orig"=original, "prReg"=motion pre-regressed)/filtering model("Loc"=locally-parameterized AR, "Glb"=globally-parameterized AR model). B) Scatterplots of parcel significance using conventional analyses (no filtering) vs. MINDy-based Filtering for each task. Yellow dots indicate significant parcels (in terms of absolute sustained effect) which also had increased effect sizes from MINDy-based Filtering, while blue dots denote significant parcels whose effect sizes were larger with conventional analyses. Teal dots denote parcels which did not exhibit a significant control-demand effect for either method. Red lines indicate the multilevel model fits to each task (linear + a bidirectional main effect of parcel significance)

sustained analyses (maximum $p = .0008$, Fig. 9A). The main effect of regional significance during multilevel modeling (see Sec. 5.5) was also greater in the proposed technique than autoregressive comparison models (local: $t(1669) = 2.49, p = .013$, global: $t(1669) = 2.36, p = .018$). However, the proposed method did not significantly magnify effect sizes over AR control models ($t = .51, .71$ for local and global, respectively). Thus, the modeling of connectivity in MINDy primarily serves to further differentiate task-implicated and non-implicated parcels. MINDy-based filtering also improved the cross-task generalizability of cognitive-control effects relative autoregressive controls at both the group-level (local ICC=.56 ± .01, global ICC=.55 ± .01 vs. MINDy-based ICC=.63 ± .01, p¡.001, 5000 bootstraps) and individual differences over the DMCC34 parcels (local: paired-$t(33) = 2.2, p = .032$, global: $t = 2.49, p = .018$).

7.3. Detection of Sustained Effects

As mentioned in the previous section, sustained effects were magnified for each task relative the two auto-regressive controls (maximum $p = .0008$, Fig. 9A). However, sustained effects during Cued-TS did not significantly

improve relative to conventional analyses (paired-$t(107) = .20, p = .84$, Fig. 9A,B) or the pre-regressed control (paired-$t(99) = 1.64, p = .10$, Fig. 9A). Sustained effects for all other tasks did significantly increase (maximum $p = 9E-5$). Collapsed across tasks, MINDy-based Filtering improved the estimation of sustained effects relative both conventional controls and autoregressive models (maximum $p = 6E-15$). Thus, the proposed technique generally increased statistical power in detecting sustained effects. MINDy-based Filtering also increased the cross-task generalizability of group-average sustained effects (MINDy=.68±.02, all controls< .6, $p < .001$, 5000 bootstraps). However, it's important to note that sustained effects are not "construct-pure" and their distribution was highly skewed (strong visual component) so we urge caution in interpreting cross-task generalizability of sustained responses (although see Sec. 5.8 for its relevance to construct-specific behavior).

7.4. Sensitivity of Sustained Effects

As with event-related analyses, we examined whether improvements in the detection of sustained effects were limited to task-implicated regions. As before, we considered bidirectional effects for sustained analyses (i.e. parcels with significant increases or decreases in sustained activity). For this reason, we slightly modified Eq. 12 to model improvements in terms of magnitude rather than a linear main effect:

$$X_{task,Parc} = \beta_{task} + \beta_0 Y_{task,Parc} + Sig_{task,Parc}(\beta_1 sign(Y_{task,Parc}) + \beta_2 Y_{task,Parc}) + \epsilon_{task,Parc}. \quad (13)$$

Note that the coefficient β_1 is now multiplied $sign(Y_{task,Parc})$. Results for sustained analysis mirrored those of the event-related analysis. As with event-related analyses, the proposed technique differentially increased effect sizes over task-implicated parcels when compared to both conventional controls ($t(1669)$=7.98, $p \approx 0$ vs. original and t=8.55, $p \approx 0$ vs. pre-regressed, Fig. 9B) and these increases were even greater relative to AR controls ($t(1669)$=10.72, $p \approx$0 vs. local and t=12.75, $p \approx$0 vs. global). As with event-related analyses, differential magnification met significance vs the original analysis ($t = 2.95, p = .0033$) but not the pre-regressed control ($t = 1.93, p = .054$) indicating that improvements largely impact a main-effect of parcel significance (i.e. increased categorical distinction between task-implicated and non-implicated parcels).

7.5. Using Alternative Techniques to Estimate Input

We also compared results using the simple, linear estimation of input signals with two more advanced techniques: either normalizing results with respect to the estimated hemodynamics or performing estimation using a nonlinear Kalman Filter. We found that HRF normalization significantly improved the detection of event-related effects in 3 of 4 tasks

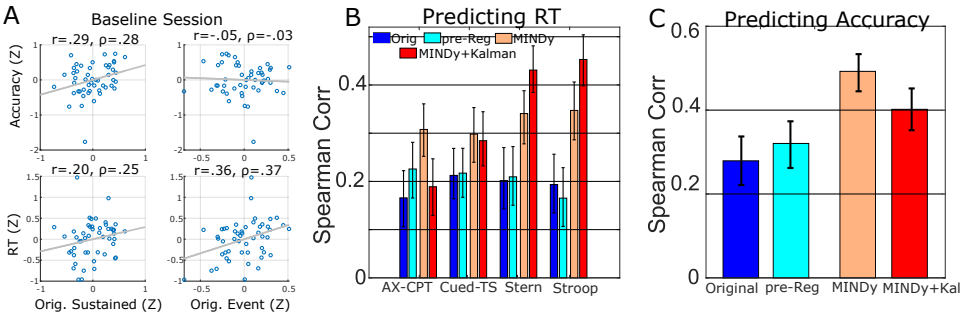

Fig 10. MINDy-based Filtering enhances prediction of individual differences in cognitive control effects. Behavioral measures correspond to the difference of high control trials minus low control trials in normalized response times (RT) and accuracy (Acc). A) Individual differences in DMCC34 sustained responses predict accuracy while event-related responses (high vs. low control contrast) predict RTs in the baseline session (collapsed across tasks). Data plotted correspond to the original (non-MINDy) analyses. B) MINDy-based Filtering enhances correlations between event-related responses and RTs in each task (collapsed across the 3 scanning sessions). C) MINDy-based Filtering also enhances the correlation between sustained responses in DMCC34 and error-rates (baseline session). Confidence intervals correspond to ±1 SD.

(maximum $p = .0014$), but worsened detection for the Stroop task (paired-$t(246) = -10.65, p \approx 0$) relative to MINDy-based Filtering without normalization. In all four tasks it significantly increased detection power relative conventional approaches ($p \approx 0$). There was no difference in event-related detection power collapsed over tasks (paired-$t(581) = .74, p = .46$). HRF normalization also slightly decreased cross-task generalizability ($ICC = .623$ vs. .635, $p < .05$, 5000 bootstraps) but did not impact generalizability of individual differences over the DMCC34 parcels (paired-$t(33) = -1.31, p = .20$).

MINDy-based Filtering with the Kalman filter improved event-related detection in 2 of 4 tasks (maximum $p = .02$) relative to the naïve filter and in 3 of 4 tasks relative to conventional analyses (AX-CPT was not significant). The naïve MINDy-based Filter had greater event-related detection in the AX-CPT and Stroop tasks ($p \approx 0$). Collapsed over tasks, detection power was greater than conventional techniques (maximum $p = 6E-10$), but less than the naïve filter (paired-$t(574) = -14.1, p \approx 0$). Detection power of sustained effects increased in 2 tasks relative to the naïve MINDy-based Filter (Cued-TS and Stroop. maximum $p = 4.7E-5$) and did not significantly differ in the other two tasks. Collapsed across tasks, the increase in detection power was significant (paired-$t(654) = 4.702, p = 3.16E-6$). MINDy-based Filtering with the Kalman Filter also increased cross-task generalizability of group-level effects ($ICC = .664 \pm .01, p < .05$, 5000 bootstraps), but decreased generalizability of individual differences in the DMCC34 parcels (paired-$t(33) = -3.05, p = .0045$) relative the naïve filter. We also performed behavioral analyses with the Kalman filter variant of MINDy based Filtering (other results for this method are reported in SI Sec. 7.5). This variant

significantly improved predictive power (mean across tasks) over both the original and pre-regressed controls ($p < .05$, 5,000 bootstraps) although the increase over simple MINDy-based Filtering was not statistically significant (SI Fig. 10 B).

7.6. Sensitivity to Motion

Lastly, we compared the sensitivity of approaches to motion artifact. For each task and scanning session we computed three motion statistics: the number of frames censored due to passing a critical value of framewise displacement, the median framewise displacement and the median DVARS statistic ([41]) for each task run and averaged over runs. We then used resampling to test the relationship between each motion variable and the group effect-size of the high-vs.-low conflict contrast and sustained effect for each task. In brief, we randomly drew 5,000 samples of 30 subjects each without replacement. We computed group-level statistics for motion and the cognitive control contrast and then tested whether the average motion or variability of motion (inter-subject) of a sample predicted the sample's group-effect (one-sample t-scores averaged over the 34 parcels). We also used the same technique for predicting the difference between methods (i.e. do improvements under our approach require low motion?). Results did not indicate a significant effect of motion for the current dataset and subject pool. The relationship between motion and the difference between methods (MINDy versus original averaged over tasks) was insignificant for event-related analyses and did not display a consistent sign (proportion of frames censored:$r = .013$, FD: $r = -.015$, DVARS: $r = -.06$). Likewise, we did not observe differential sensitivity to motion in the sustained effects (frames censored: $r = 0$, FD: $r = .002$, DVARS: $r = -.002$). Thus, the degree to which MINDy-based Filtering improves upon conventional methods is not influenced by motion within reasonable bounds.

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.

References

[1] M. F. Singh, T. S. Braver, M. W. Cole, S. Ching, Estimation and validation of individualized dynamic brain models with resting state fmri,

Table 1. Attributes of the DMCC34 parcels. Indices are for the Schaefer 400-region, 7-network parcellation [31]). Coordinates (X,Y,Z) refer to MNI centroids.

Parcel Number	Hemisphere	Network	ROI	X	Y	Z
22	Left	Visual	22	-19	-65	7
77	Left	Dorsal Attention	Post_9	-33	-46	41
78	Left	Dorsal Attention	Post_10	-29	-58	50
86	Left	Dorsal Attention	FEF_1	-40	-3	51
87	Left	Dorsal Attention	FEF_2	-25	-1	55
91	Left	Dorsal Attention	PrCv_2	-50	3	38
93	Left	Salience/Ventral Attention	ParOper_2	-58	-44	27
99	Left	Salience/Ventral Attention	FrOperIns_3	-33	25	-1
101	Left	Salience/Ventral Attention	FrOperIns_5	-33	19	8
103	Left	Salience/Ventral Attention	FrOperIns_7	-43	12	2
105	Left	Salience/Ventral Attention	FrOperIns_9	-52	9	13
107	Left	Salience/Ventral Attention	Med_1	-6	22	31
110	Left	Salience/Ventral Attention	Med_4	-5	9	48
127	Left	Control	Par_1	-29	-74	42
130	Left	Control	Par_4	-35	-62	48
139	Left	Control	PFCl_5	-42	38	22
140	Left	Control	PFCl_6	-45	20	27
144	Left	Control	pCun_1	-9	-77	45
148	Left	Control	PFCmp_1	-4	28	47
172	Left	Default	PFC_7	-48	28	0
185	Left	Default	PFC_10	-53	19	11
185	Left	Default	PFC_20	-42	7	48
189	Left	Default	PFC_24	-6	10	65
219	Right	Visual	19	9	-74	9
301	Right	Salience/Ventral Attention	PrC_1	51	3	41
303	Right	Salience/Ventral Attention	FrOperIns_2	41	8	-3
306	Right	Salience/Ventral Attention	FrOperIns_5	37	23	5
314	Right	Salience/Ventral Attention	Med_4	6	11	58
340	Right	Control	PFCv_1	34	21	-8
346	Right	Control	PFCl_6	50	30	18
347	Right	Control	PFCl_7	48	18	23
349	Right	Control	PFCl_9	47	29	28
350	Right	Control	PFCl_10	39	11	34
353	Right	Control	PFCl_13	43	7	51

NeuroImage 221 (2020) 117046. doi:https://doi.org/10.1016/j.neuroimage.2020.117046.

[2] B. Biswal, F. Z. Yetkin, V. M. Haughton, J. S. Hyde, Functional connectivity in the motor cortex of resting human brain using echo-planar mri, Magnetic Resonance in Medicine 34 (4) (1995) 537–541. doi:10.1002/mrm.1910340409.

[3] E. A. Ashley, The Precision Medicine Initiative: A New National Effort, JAMA 313 (21) (2015) 2119–2120. doi:10.1001/jama.2015.3595.

[4] B. M. Psaty, O. M. Dekkers, R. S. Cooper, Comparison of 2 Treatment Models: Precision Medicine and Preventive Medicine, JAMA 320 (8) (2018) 751–752. doi:10.1001/jama.2018.8377.

[5] T. D. Satterthwaite, C. H. Xia, D. S. Bassett, Personalized neuroscience: Common and individual-specific features in functional brain networks, Neuron 98 (2) (2018) 243 – 245. doi:10.1016/j.neuron.2018.04.007.

[6] M. Mennes, C. Kelly, X.-N. Zuo, A. D. Martino, B. B. Biswal, F. X. Castellanos, M. P. Milham, Inter-individual differences in resting-state functional connectivity predict task-induced bold activity, NeuroImage 50 (4) (2010) 1690 – 1701. doi:doi.org/10.1016/j.neuroimage.2010.01.002.

[7] I. Tavor, O. P. Jones, R. B. Mars, S. M. Smith, T. E. Behrens, S. Jbabdi, Task-free mri predicts individual differences in brain activity during task performance, Science 352 (6282) (2016) 216–220. doi:10.1126/science/aad8127.

[8] M. W. Cole, T. Ito, D. S. Bassett, D. H. Schultz, Activity flow over resting-state networks shapes cognitive task activations, Nature Neuroscience 19 (12) (2016) 1718–1726. doi:10.1038/nn.4406.

[9] E. M. Gordon, T. O. Laumann, A. W. Gilmore, D. J. Newbold, D. J. Greene, J. J. Berg, M. Ortega, C. Hoyt-Drazen, C. Gratton, H. Sun, J. M. Hampton, R. S. Coalson, A. L. Nguyen, K. B. McDermott, J. S. Shimony, A. Z. Snyder, B. L. Schlaggar, S. E. Petersen, S. M. Nelson, N. U. Dosenbach, Precision functional mapping of individual human brains, Neuron 95 (4) (2017) 791 – 807.e7. doi:10.1016/j.neuron.2017.07.011.

[10] B. J. He, Spontaneous and task-evoked brain activity negatively interact, Journal of Neuroscience 33 (11) (2013) 4672–4682. doi:10.1523/JNEUROSCI.2922-12.2013.

[11] M. D. Fox, A. Z. Snyder, J. M. Zacks, M. E. Raichle, Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses, Nature Neuroscience 9 (2006) 23–25. doi:10.1038/nn1616.

[12] A. Ponce-Alvarez, B. J. He, P. Hagmann, G. Deco, Task-driven activity reduces the cortical activity space of the brain: experiment and whole-brain modeling, PLoS computational biology 11 (8) (2015) e1004445. doi:10.1371/journal.pcbi.1004445.

[13] M. Demirtaş, J. B. Burt, M. Helmer, J. L. Ji, B. D. Adkinson, M. F. Glasser, D. C. Van Essen, S. N. Sotiropoulos, A. Anticevic, J. D. Murray, Hierarchical heterogeneity across human cortex shapes large-scale neural dynamics, Neuron 101 (6) (2019) 1181–1194. doi:10.1016/j.neuron.2019.01.017.

[14] P. Wang, R. Kong, X. Kong, R. Liégeois, C. Orban, G. Deco, M. P. van den Heuvel, B. Thomas Yeo, Inversion of a large-scale circuit model reveals a cortical hierarchy in the dynamic resting human brain, Science Advances 5 (1) (2019). doi:10.1126/sciadv.aat7854.

[15] B. J. He, J. M. Zempel, Average is optimal: An inverted-u relationship between trial-to-trial brain activity and behavioral performance, PLOS Computational Biology 9 (11) (2013) 1–14. doi:10.1371/journal.pcbi.1003348.

[16] K. E. Stephan, L. Kasper, L. M. Harrison, J. Daunizeau, H. E. [den Ouden], M. Breakspear, K. J. Friston, Nonlinear dynamic causal models for fmri, NeuroImage 42 (2) (2008) 649 – 662. doi:10.1016/j.neuroimage.2008.04.262.

[17] M. Singh, A. Wang, T. Braver, S. Ching, Scalable surrogate deconvolution for identification of partially-observable systems and brain modeling, Journal of Neural Engineering (2020).
URL http://iopscience.iop.org/10.1088/1741-2552/aba07d

[18] M. D. Fox, A. Z. Snyder, J. L. Vincent, M. E. Raichle, Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior, Neuron 56 (1) (2007) 171 – 184. doi:10.1016/j.neuron.2007.08.023.

[19] K. Friston, L. Harrison, W. Penny, Dynamic causal modelling, NeuroImage 19 (4) (2003) 1273 – 1302. doi:10.1016/S1053-8119(03)00202-7.

[20] A. Razi, M. L. Seghier, Y. Zhou, P. McColgan, Large-scale dcms for resting-state fmri, Network Neuroscience 1 (3) (2017). doi:10.1162/NETN_a_00015.

[21] S. Frässle, E. I. Lomakina, A. Razi, K. J. Friston, J. M. Buhmann, K. E. Stephan, Regression dcm for fmri, Neuroimage 155 (2017) 406–421.

[22] K. Friston, K. H. Preller, C. Mathys, H. Cagnan, J. Heinzle, A. Razi, P. Zeidman, Dynamic causal modelling revisited, NeuroImage 199 (2019) 730 – 744. doi:10.1016/j.neuroimage.2017.02.045.

[23] N. Weiner, Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications, MIT Press, New York City, 1949.

[24] T. S. Braver, A. Kizhner, R. Tang, M. C. Freund, J. A. Etzel, The dual mechanisms of cognitive control (dmcc) project, bioRxiv (2020). arXiv:https://www.biorxiv.org/content/early/2020/09/20/2020.09.18.304402.full.pdf, doi:10.1101/2020.09.18.304402.

[25] J. Cohen, D. Barch, C. Carter, D. Servan-Schreiber, Schizophrenic deficits in the processing of context: Converging evidence from three theoretically motivated cognitive tasks, Journal of Abnormal Psychology 108 (1999) 120–133.

[26] S. Sternberg, High-speed scanning in human memory, Science 153 (1966) 652–654. doi:10.1126/science.153.3736.652.

[27] J. Stroop, Studies of interference in serial verbal reactions, Journal of Experimental Psychology 18 (1935) 643–662.

[28] J. Bugg, T. Braver, Proactive control of irrelevant task rules during cued task switching, Psychological Research 80 (2016) 860–876. doi:10.1007/s00426-015-0686-5.

[29] O. Esteban, C. Markiewicz, R. e. a. Blair, fmriprep: a robust preprocessing pipeline for functional mri, Nature Methods 16 (2019) 111–116. doi:https://doi.org/10.1038/s41592-018-0235-4.

[30] O. Esteban, R. Ciric, K. e. a. Finc, Analysis of task-based functional mri data preprocessed with fmriprep, Nature Protocols 15 (2020) 2186–2202. doi:https://doi.org/10.1038/s41596-020-0327-3.

[31] A. Schaefer, R. Kong, E. M. Gordon, T. O. Laumann, X.-N. Zuo, A. J. Holmes, S. B. Eickhoff, B. T. Yeo, Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri, Cerebral Cortex (2017) 1–20doi:10.1093/cercor/bhx179.

[32] B. Fischl, Freesurfer, NeuroImage 62 (2) (2012) 774 – 781. doi:10.1016/j.neuroimage.2012.01.021.

[33] G. H. Glover, Deconvolution of impulse response in event-related bold fmri1, NeuroImage 9 (4) (1999) 416 – 429. doi:10.1006/nimg.1998.0419.

[34] C. Goutte, F. A. Nielsen, K. H. Hansen, Modeling the hemodynamic response in fmri using smooth fir filters, IEEE Transactions on Medical Imaging 19 (12) (2000) 1188–1201. doi:10.1109/42.897811.

[35] J. Ollinger, G. Shulman, M. Corbetta, Separating processes within a trial in event-related functional mri: I. the method, NeuroImage 13 (1) (2001) 210 – 217. doi:10.1006/nimg.2000.0710.

[36] R. W. Cox, Afni: Software for analysis and visualization of functional magnetic resonance neuroimages, Computers and Biomedical Research 29 (1996) 162–173.

[37] K. M. Visscher, F. M. Miezin, J. E. Kelly, R. L. Buckner, D. I. Donaldson, M. P. McAvoy, V. M. Bhalodia, S. E. Petersen, Mixed blocked/event-related designs separate transient and sustained activity in fmri, Neuroimage 19 (4) (2003) 1694–1708. doi:10.1016/s1053-8119(03)00178-2.

[38] J. W. D. Steven E Petersen 1, The mixed block/event-related design, Neuroimage 62 (2) (2012) 1177–1184. doi:10.1016/j.neuroimage.2011.09.084.

[39] T. Lefebvre, H. Bruyninckx, J. De Schutter, Kalman filters for non-linear systems: a comparison of performance, International journal of Control 77 (7) (2004) 639–653.

[40] P. Shrout, J. Fleiss, Intraclass correlations: uses in assessing rater reliability, Psychological bulletin 86 (2) (1979) 420–428.

[41] J. D. Power, K. A. Barnes, A. Z. Snyder, B. L. Schlaggar, S. E. Petersen, Spurious but systematic correlations in functional connectivity mri networks arise from subject motion, NeuroImage 59 (3) (2012) 2142 – 2154. doi:10.1016/j.neuroimage.2011.10.018.

[42] B. T. Thomas Yeo, F. M. Krienen, J. Sepulcre, M. R. Sabuncu, D. Lashkari, M. Hollinshead, J. L. Roffman, J. W. Smoller, L. Zöllei, J. R. Polimeni, B. Fischl, H. Liu, R. L. Buckner, The organization of the human cerebral cortex estimated by intrinsic functional connectivity, Journal of Neurophysiology 106 (3) (2011) 1125–1165. doi:10.1152/jn.00338.2011.

[43] B. Biswal, E. A. Deyoe, J. S. Hyde, Reduction of physiological fluctuations in fmri using digital filters, Magnetic Resonance in Medicine 35 (1) (1996) 107–113. doi:10.1002/mrm.1910350114.

[44] O. Friman, M. Borga, P. Lundberg, H. Knutsson, Detection and detrending in fmri data analysis, NeuroImage 22 (2) (2004) 645 – 655. doi:https://doi.org/10.1016/j.neuroimage.2004.01.033.

[45] C. E. Davey, D. B. Grayden, G. F. Egan, L. A. Johnston, Filtering induces correlation in fmri resting state data, NeuroImage 64 (2013) 728 – 740. doi:https://doi.org/10.1016/j.neuroimage.2012.08.022.

[46] S. Särkkä, A. Solin, A. Nummenmaa, A. Vehtari, T. Auranen, S. Vanni, F.-H. Lin, Dynamic retrospective filtering of physiological noise in bold fmri: Drifter, NeuroImage 60 (2) (2012) 1517 – 1527. doi:https://doi.org/10.1016/j.neuroimage.2012.01.067.

[47] T. D. Satterthwaite, M. A. Elliott, R. T. Gerraty, K. Ruparel, J. Loughead, M. E. Calkins, S. B. Eickhoff, H. Hakonarson, R. C. Gur, R. E. Gur, D. H. Wolf, An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data, NeuroImage 64 (2013) 240 – 256. doi:https://doi.org/10.1016/j.neuroimage.2012.08.052.

[48] M. W. Cole, D. S. Bassett, J. D. Power, T. S. Braver, S. E. Petersen, Intrinsic and task-evoked network architectures of the human brain, Neuron 83 (1) (2014) 238 – 251. doi:10.1016/j.neuron.2014.05.014.

[49] G. D. Logan, W. B. Cowan, On the ability to inhibit thought and action: A theory of an act of control, Psychological Review 91 (3) (1984) 295–327. doi:10.1037/0033-295X.91.3.295.

[50] M. Usher, J. L. McClelland, The time course of perceptual choice: The leaky, competing accumulator model, Psychological Review 108 (3) (2001) 550–592. doi:10.1037/0033-295X.108.3.550.

[51] N. P. Rougier, D. C. Noelle, T. S. Braver, J. D. Cohen, R. C. O'Reilly, Prefrontal cortex and flexible cognitive control: Rules without symbols, Proceedings of the National Academy of Sciences 102 (20) (2005) 7338–7343. doi:10.1073/pnas.0502455102.

[52] F. Verbruggen, G. D. Logan, Models of response inhibition in the stop-signal and stop-change paradigms, Neuroscience and Biobehavioral Reviews 33 (5) (2009) 647 – 661. doi:doi.org/10.1016/j.neubiorev.2008.08.014.

[53] T. T. Rogers, J. L. McClelland, Parallel distributed processing at 25: Further explorations in the microstructure of cognition, Cognitive Science 38 (6) (2014) 1024–1077.

[54] H. Fields, S. Anderson, Evidence that raphe-spinal neurons mediate opiate and midbrain stimulation-produced analgesias, Pain 5 (4) (1978) 333–349. doi:10.1016/0304-3959(78)90002-7.

[55] A. London, I. Benhar, M. Schwartz, The retina as a window to the brain—from eye research to cns disorders, Nature Reviews Neurology 9 (2013) 44–53. doi:10.1038/nrneurol.2012.227.

[56] T. J. Buschman, E. L. Denovellis, C. Diogo, D. Bullock, E. K. Miller, Synchronous oscillatory neural ensembles for rules in the prefrontal cortex, Neuron 76 (4) (2012) 838 – 846. doi:10.1016/j.neuron.2012.09.029.

[57] T. Akam, D. Kullmann, Oscillatory multiplexing of population codes for selective communication in the mammalian brain, Nature Reviews Neuroscience 15 (2014) 111–122. doi:10.1038/nrn3668.

[58] E. Smith, G. Horga, M. e. a. Yates, Widespread temporal coding of cognitive control in the human prefrontal cortex, Nature Neuroscience 22 (2019) 1883–1891. doi:10.1038/s41593-019-0494-0.