



Ground-truth “resting-state” signal provides data-driven estimation and correction for scanner distortion of fMRI time-series dynamics

Rajat Kumar^a, Liang Tan^b, Alan Kriegstein^b, Andrew Lithen^a, Jonathan R. Polimeni^{c,d},
Lilianne R. Mujica-Parodi^{a,c,e}, Helmut H. Strey^{a,e,*}

^a Department of Biomedical Engineering, Stony Brook University, School of Medicine, Stony Brook, NY 11794-5281, USA

^b ALA Scientific Instruments, Inc., Farmingdale, NY, USA

^c Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Harvard Medical School, Massachusetts General Hospital, Charlestown, MA, USA

^d Harvard-Massachusetts Institute of Technology Division of Health Sciences and Technology, Cambridge, MA, USA

^e Laufer Center for Physical and Quantitative Biology, Stony Brook University, Stony Brook, NY, USA

ARTICLE INFO

Keywords:

Dynamic phantom
Scanner instability
Multi-site
Marchenko-Pastur distribution
Dynamic fidelity

ABSTRACT

The fMRI community has made great strides in decoupling neuronal activity from other physiologically induced T_2^* changes, using sensors that provide a ground-truth with respect to cardiac, respiratory, and head movement dynamics. However, blood oxygenation level-dependent (BOLD) time-series dynamics are also confounded by scanner artifacts, in complex ways that can vary not only between scanners but even, for the same scanner, between sessions. Unfortunately, the lack of an equivalent ground truth for BOLD time-series has thus far stymied the development of reliable methods for identification and removal of scanner-induced noise, a problem that we have previously shown to severely impact detection sensitivity of resting-state brain networks. To address this problem, we first designed and built a phantom capable of providing dynamic signals equivalent to that of the resting-state brain. Using the dynamic phantom, we then compared the ground-truth time-series with its measured fMRI data. Using these, we introduce data-quality metrics: Standardized Signal-to-Noise Ratio (ST-SNR) and Dynamic Fidelity that, unlike currently used measures such as temporal SNR (tSNR), can be directly compared across scanners. Dynamic phantom data acquired from four “best-case” scenarios: high-performance scanners with MR-physicist-optimized acquisition protocols, still showed scanner instability/multiplicative noise contributions of about 6–18% of the total noise. We further measured strong non-linearity in the fMRI response for all scanners, ranging between 8–19% of total voxels. To correct scanner distortion of fMRI time-series dynamics at a single-subject level, we trained a convolutional neural network (CNN) on paired sets of measured vs. ground-truth data. The CNN learned the unique features of each session’s noise, providing a customized temporal filter. Tests on dynamic phantom time-series showed a 4- to 7-fold increase in ST-SNR and about 40–70% increase in Dynamic Fidelity after denoising, with CNN denoising outperforming both the temporal bandpass filtering and denoising using Marchenko-Pastur principal component analysis. Critically, we observed that the CNN temporal denoising pushes ST-SNR to a regime where signal power is higher than that of noise ($ST-SNR > 1$). Denoising human-data with ground-truth-trained CNN, in turn, showed markedly increased detection sensitivity of resting-state networks. These were visible even at the level of the single-subject, as required for clinical applications of fMRI.

1. Introduction

Large-scale investments in the identification of fMRI-derived biomarkers for brain-based disorders are a testament to the anticipated promise of fMRI as a neurodiagnostic tool. Yet even once clinical neuroscience establishes reliable biomarkers, a critical rate-limiting factor in the use of fMRI in clinical practice will be fMRI’s poor signal/noise profile for single-subject level analyses. The task-free, “resting-state”

paradigms most likely to be utilized in a clinical setting (because of their limited reliance on patient training, engagement, and compliance) only exacerbate this problem. Task-based designs, in principle, clearly delineate between activation in response to a task (signal) and activation during baseline (noise). However, task-free paradigms, by definition, lack the experimental manipulation that would typically be used to distinguish between *fluctuations of interest* (signal) from *fluctuations of nuisance* (noise) (DeDora et al., 2016). Without a principled way to

* Corresponding author at: Department of Biomedical Engineering, Stony Brook University, School of Medicine, Stony Brook, NY 11794-5281, USA
E-mail addresses: Lilianne.Strey@stonybrook.edu (L.R. Mujica-Parodi), Helmut.Strey@stonybrook.edu (H.H. Strey).

distinguish between signal and noise, we lack the feedback necessary to optimize for one while removing the other, thereby limiting our ability to achieve the kind of advances in detection sensitivity required to enhance fMRI's utility in evaluating the single patient.

fMRI's signal is conventionally derived from the *blood oxygenation level-dependent* (BOLD) contrast. This activity represents regional time-varying changes in the concentration of deoxygenated hemoglobin, following neural-activity induced by exogenous stimuli or spontaneous fluctuations of the resting state. These time-varying changes reflect changes in apparent transverse relaxation time T_2^* , an MR parameter sensitive to levels of deoxyhemoglobin, and hence responsible for the observed BOLD contrast. Ideally, the value measured at each voxel at a given time point should only change in response to T_2^* changes driven by neural activity (fluctuations of interest, signal). However, in practice, the measurement is dependent on a complex interaction between acquisition parameters (flip-angle: α , echo-time: TE, repetition-time: TR), MR parameters (longitudinal relaxation time T_1 , apparent transverse relaxation time T_2^* , proton density within a voxel) and background noise (Lauterbur, 2000). Change in any of these parameters introduces variance (fluctuations of the nuisance, noise) in the observed voxel time-series. The difficulty of maintaining fidelity to actual (neuronal) time-series dynamics is made even more acute by the fact that BOLD contrast constitutes only a small fraction (typically, less than 5%) of the total measured signal.

Fluctuations of nuisance in the fMRI time-series originate from two sources: the individual being scanned (physiological noise, due primarily to cardiac, respiratory, and motion effects) as well as the scanner itself. Physiological processes like respiration or cardiac pulsations can cause changes in blood flow (affecting T_1 and T_2^*), and thus temporal variations in magnetization that might artifactually appear to be a BOLD effect. Subject head motion causes relative displacement of voxels leading to temporally correlated non-stationary noise and can induce spurious correlations in the resting-state analysis (Power et al., 2012). As significant as these artifacts are, the fact that cardiac, respiratory, and motion variables permit external measurements (e.g., ECG for heart rate) have permitted the field to develop an impressive array of well-validated methods with which to both identify and mitigate their influence. Examples of strategies for targeting physiological and motion confounds include: selecting acquisition parameters designed to permit thermal noise to dominate physiological noise (Wald and Polimeni, 2017); techniques to address breathing-related field fluctuations both prospectively (Duerst et al., 2015) and at image reconstruction stage (Bollmann et al., 2017); use of simultaneously recorded measurement of heart-rate, respiration, and motion to retrospectively remove physiological confounds (Caballero-Gaudes and Reynolds, 2017); and motion-correction implemented prospectively (Zaitsev et al., 2017) or retrospectively through registration.

In contrast, the lack of a ground truth for fMRI time-series has not permitted the same strategies for identification and removal of scanner-induced noise, which can vary not only between scanners of the same make and model, but even within the same scanner during different sessions. These fluctuations of nuisance originate from imperfections of the instrumentation and the electromagnetic fields used for the measurement and are normally referred to as “scanner instability.” This nomenclature is, itself, potentially misleading, since detection-sensitivity of resting-state networks requires simultaneously amplifying fluctuations of interest while suppressing fluctuations of nuisance. Indeed, we have previously shown that typical methods that focus entirely on suppressing fluctuations (optimizing solely for scanner “stability”), such as temporal signal/noise (tSNR), actually *deoptimize* detection-sensitivity of resting-state networks, because the damped fluctuations include not only suppressed noise but also suppressed signal (DeDora et al., 2016).

Different approaches tackle the problem of minimizing scanner artifacts based upon models of MR-physics. Such methods include reducing the effects of eddy currents by the use of actively shielded gradients and pre-emphasis filters, the use of navigators and calibration echoes,

or NMR probes (Kasper et al., 2015) that provide concurrent field monitoring with correction during image reconstruction. Yet modeling-based approaches, while valuable in their own right in terms of contributing to our understanding, can fall short as a practical tool for optimizing resting-state signal/noise (SNR). The reason for this is that they tend to oversimplify processes that, in an actual testing environment, are fundamentally complex—involving multiple factors, both known and unknown, which interact with one another in nonlinear and nonstationary ways. For example, scanner instabilities may be caused by variation in flip angle over time, imperfections in gradient system, heating, time-varying eddy current effects, or gain changes in transmit and receive chains (Greve et al., 2013; Liu, 2016). Time-varying gradients in fast imaging methods, such as interleaved echo-planar imaging (EPI), require high-gradient amplitudes and slew-rates, pushing the scanner to its limits and causes image artifacts due to k-space trajectory deviations. Inhomogeneity in B_0 field and perturbations in gradient field cause eddy currents, ghosting, geometric distortions, errors in phase encoding leading to voxel displacement, gain-drifts, and other distortions (Jezzard and Clare, 1999). While scanner instability is multiplicative, the impact of thermal/background noise on fMRI time-series is additive and can arise due to a random process like Brownian motion of ions in MR electronics or the human subject, external RF noise sources in the scanner room, or RF spikes due to intermittent contact between metallic components (Greve et al., 2013; Liu, 2016).

In a clinical setting involving decision-making for a single patient, the impact of errors that fluctuate over time and are signal dependent cannot be remedied by increasing sample size, under the assumption that signal amplifies while noise cancels. Longitudinal comparison of scans acquired pre and post treatment cannot be interpreted if both the subject and scanner are changing over time (for example, in using resting-state fMRI in pre-surgical localization, surgical planning in epilepsy, and identifying subjects with Alzheimer's Disease (Lee et al., 2013)). Moreover, biomarkers used at one site may be difficult to compare across other sites. Even in the research domain, recent years have seen a tremendous increase in efforts in pooling fMRI data for increasing sample size, enhancing statistical power for detecting subtle effects, including diverse populations and disease etiologies (Van Horn and Toga, 2009), either via multi-site studies or data-sharing initiatives. Combining data from multiple sites presents an unavoidable challenge in the form of scanner-induced inter-site variability due to differences in field strength, imaging parameters, image reconstruction, or scanner manufacturer (Glover et al., 2012) and can lead to systematic confounds in time-series data. In one recent example (Friedman et al., 2008), between-site reliability showed median intra-class correlation of just $r=0.22$.

In summary, efforts to make the application of resting-state fMRI clinically useful must necessarily address SNR from the perspective of not only physiological, but scanner, artifact—and in ways that make sense given the ubiquity of task-free designs. While efforts to mitigate physiological artifact can and have benefited from external measurements (Caballero-Gaudes and Reynolds, 2017), until recently such a strategy has not been available for scanner artifact. Static phantoms optimize purely for general stability (Friedman and Glover, 2006), thereby suppressing the fluctuations responsible for resting-state signal. Moreover, the brain (non-static but, by definition, the unknown variable) likewise cannot serve as a calibration device. Finally, physics-based models cannot, in principle, approximate the impact of complex nonstationary distortion on time-series without empirical measurement of that distortion. To address these issues, we approached the problem from the perspective of creating a “brain-like” calibration device, capable of producing a dynamic ground-truth input signal similar to a typical resting-state time-series. Because such a device would provide a ground truth for both fluctuations of interest (signal) as well as fluctuations of nuisance (noise), it could permit optimization for signal-to-noise, rather than simply stability. Because of the consequent ability to obtain, and therefore compare, time-series distortion between true and measured

time-series, we could develop a purely data-driven—rather than modeled—distortion correction. Doing so would potentially permit cleaning data of scanner-induced artifact while remaining agnostic with respect to the diversity of known and unknown sources of distortion and their behavior over time.

2. Results

2.1. We designed and engineered a commercial-grade dynamic phantom capable of producing brain-like dynamic signals.

Our previous work (Rădulescu and Mujica-Parodi, 2014; Mujica-Parodi et al., 2017) and those of others (Ciuciu et al., 2012) shows that healthy resting-state fMRI signals follow $1/f$ (pink noise) frequency spectra; therefore, our pseudo-brain “input” signal was engineered to achieve equivalent dynamics (custom dynamics can also be easily programmed). To create a dynamic signal, our phantom (Fig. 1A) uses difference in agarose gel concentration across voxels; the phantom, when rotated in-plane across a voxel during the data acquisition, produces a changing T_2^* signal. Rotations occur at the start of each TR of a scan and are limited to around 250 milliseconds.

The phantom consists of three distinct parts: a) an agarose gel cylinder assembly, having two concentric cylinders; b) a control unit providing control logic for rotation of the inner cylinder; and, c) an air motor assembly with a gearbox and an optical encoder for position tracking. Within the agarose gel cylinder assembly, the inner cylinder rotates during the scan and is coupled to the air motor and the optical encoder, while the outer cylinder contains a reference gel and remains static. The outer cylinder's reference agarose gel is made at 2.2% concentration by weight, whereas the inner cylinder contains two different gel concentrations at 2.2% and 2.3% by weight, split into four quadrants in a configuration as shown in Fig. 1B. Within each quadrant, a variation in T_2^* values exist across voxels because of imperfect agarose network formation, chemical heterogeneity, and polydispersity of gel networks (Djabourov et al., 1989). The control unit for driving the phantom uses a feedback control strategy with control logic implemented in PSoC microcontroller, feedback sensing via an optical encoder, and actuation through solenoid valves. The control unit contains some other custom circuitry for fast valve response time (spike-up voltage circuit), touch-screen user-interface running on raspberry-pi, and UART communication between the raspberry-pi and the PSoC microcontroller. The phantom is MR-compatible (agarose gel cylinder assembly and air motor assembly) and uses polycarbonate (body), delrin (air motor), glass-nylon (ball bearings), and G11 garolite (motor shaft) in construction. The control unit containing electronics and pneumatic compressor for driving the air motor stays outside in the MR control room.

2.2. Using ground truth brain-like dynamic signals, we quantified a Standardized Signal-to-Noise Ratio (ST-SNR) and Dynamic Fidelity; these demonstrated wide variance across scanners, even for the “best case scenario” of high-performance scanners utilizing acquisition parameters individually optimized by a highly experienced MR physicist.

While the definition of signal-to-noise ratio (SNR) is well defined across the engineering domain, use of the term within the fMRI field has colloquially co-opted its definition in ways that can dilute its meaning and utility. Currently in fMRI, multiple definitions and variants for computing SNR exist (Welvaert and Rosseel, 2013), leading to difficulty in interpreting and comparing SNR values. Normally used to optimize for scanner stability with the use of a static phantom, *temporal SNR* (tSNR) is defined as the ratio of mean signal to standard deviation of a time-series. However, for reasons described above, optimizing for tSNR (i.e., solely for stability) will suppress not only the fluctuations responsible for noise but also the fluctuations responsible for resting-state signal, effectively de-optimizing for detection of resting-state networks (DeDora et al., 2016). Furthermore, mean-signal in tSNR calculation is

highly dependent on acquisition parameters, making the interpretation for comparison difficult. For example: tSNR has been reported across two orders of magnitude (e.g., between 4.42 and 280 for a recent review of studies (Welvaert and Rosseel, 2013)). To address both issues, we quantified the accuracy with which fMRI time-series follow the true signal using two data-quality metrics: *Standardized Signal-to-Noise Ratio (ST-SNR)* and *Dynamic Fidelity*. “ST-SNR” is defined as the ratio of signal power and the background noise power and is calculated accordingly, where power is the sum of the absolute squares of time-domain samples divided by the time-series length. We define “Dynamic Fidelity” as the accuracy with which an MR scanner tracks changes in the input signal and calculate it as the Pearson correlation coefficient between the ground-truth signal and fMRI output.

The programmed rotation of the dynamic phantom, along with the optical encoder feedback, provides a mechanism for rotation control and sensing. The phantom tracks the programmed rotation at an accuracy of 0.2° . With the rotation generating voxel-wise time-series, the feedback sensing provides data on the actual rotation that occurs. This feedback data enables calculation of the ground-truth time-series and the noise estimate for each voxel, as shown in Fig. 2, for quantifying ST-SNR and Dynamic Fidelity. In Table 1, we show both ST-SNR and Dynamic Fidelity for four scanners, showing the potential for wide variance across scanners, even for a “best case scenario” of high-performance scanners utilizing acquisition parameters individually optimized by a highly experienced MR physicist. Importantly for multi-site or longitudinal applications, these two metrics (ST-SNR and Dynamic Fidelity) provide a direct assessment and comparison of data-quality over different scanners, as well as the same scanner over time. As ST-SNR and Dynamic Fidelity have standardized and interpretable range of values, the direct comparison of these metrics longitudinally or across scanners becomes possible. For example, for the same make and model of a scanner (the two Siemens PRISMA scanners, described in Table 1) having equivalent voxel-size, the ST-SNR observed is markedly different. Inspecting further, while one may attribute this difference to the different head-coil arrays used between the two scanners (Table 3), the comparison of ST-SNR with 3T SKYRA (Table 1) at the same site with equivalent voxel-size and head-coil suggests otherwise: that the Site 2 PRISMA scanner is an outlier.

2.3. Using the dynamic phantom generated ground-truth, we quantified the ratio of scanner instability to background noise in fMRI time-series, thereby identifying multiplicative versus thermal noise components.

We analyzed time-series of the noise (residual time-series as calculated above, refer to Fig. 2), using power spectral density plots, to identify spectral-features arising from scanner artifacts. Fig. 3 illustrates the mean power spectral density across all voxels for the ground-truth and the estimated noise time-series. Each voxel time-series' power spectral density was normalized by its maximum power before calculating mean at each frequency bin across all voxels. The power spectral density of the noise closely matches that of the ground-truth signal, indicating the presence of a multiplicative noise (scanner instability) component alongside thermal/background noise.

Multiplicative noise modulates the MR signal, is known to exhibit some temporal and spatial correlation (Greve et al., 2013), and cannot be removed using smoothing or frequency-based temporal filtering. The presence of multiplicative noise diminishes the advantages offered by hardware improvements (increase in signal to thermal noise ratio with higher field strength and more sensitive head-coil arrays) and can exacerbate the false-positives problem (Eklund et al., 2016), alongside spurious correlations and poor reproducibility of functional connectivity. Band-limited programmed rotation of our phantom produces a band-limited ground-truth signal, and thus the associated multiplicative noise can be directly observed in this narrow band—see Fig. 3 (in the 0–0.1 Hz range). At frequencies higher than 0.1 Hz where the ground-truth signal is absent, scanner noise shows a flat spectrum or white-noise be-




Fig. 1. A. Isometric view of the Dynamic Phantom. The cylindrical head is the agarose gel cylinder assembly, which is coupled to a pneumatic motor and an optical encoder on the other end. All components remain intact via fastening to an outer frame and go inside the MR scanner with the cylindrical head placed inside the head-coil. The black box shown is the control unit, which interfaces with the optical encoder, pneumatic input from an air compressor and the pneumatic motor. **B. Distribution of T_2^* values across voxels in four quadrants at 3T (Site 1).** The agarose gel is prepared using the recipe provided by Friedman et al. (Friedman and Glover, 2006). Even though the agarose gel is prepared only at 2.2% and 2.3% concentration, the heterogeneity in T_2^* values can be attributed to imperfect agarose network formation, chemical heterogeneity, and polydispersity of gel networks (Djabourov et al., 1989). **C. Feedback control system for rotating the inner cylinder.** At each trigger from the MR scanner, the PSoc controller compares the current position $F(t)$ with the programmed target position $R(t)$ and opens the solenoid valve proportionally to the magnitude of the error signal $E(t)$ to actuate the pneumatic motor. Here, $U(t)$ is the actuating signal, and $M(t)$ is the manipulated variable. The system uses no braking mechanism, and accurate positioning is achieved through a predetermined linear relationship established between open-state time for solenoid valve and the corresponding rotation achieved at a given pneumatic pressure.

havior (thermal noise). This multiplicative noise behavior is further corroborated by a linear scaling of noise power (logarithmic scale) with an increase in signal (ground-truth) standard deviation. We observed a moderate correlation between noise power and signal standard deviation for all scanners (Site 1: PRISMA- $r=0.35$, Site 2: PRISMA- $r=0.32$, SKYRA- $r=0.33$, and MAGNETOM- $r=0.35$).

Using the ground-truth dynamic signal and the measured fMRI output, we quantified the ratio of multiplicative noise (scanner instability) to thermal/background noise using a probabilistic description of the two noise-sources. Scanner instability is signal-dependent and thus proportional to the signal intensity, while thermal noise is independent of the

MR signal. Background noise and scanner-instability are temporally independent, and therefore, their variances add. With σ_T as the standard deviation of the thermal noise and β as the proportionality constant for the multiplicative noise, we can write:

$$\sigma_{fMRI}^2 = \sigma_{GT}^2 + \sigma_T^2 + \beta^2 \sigma_{GT}^2 = \sigma_{GT}^2 + \sigma_{noise}^2, \quad (1)$$

where σ_{fMRI}^2 and σ_{GT}^2 , are the variances of the observed fMRI output and the ground-truth, respectively. The model for the probability of observing the measured signal, given ground-truth Y_{GT} and noise parameters can then be written as:

$$P(Y_{fMRI} | Y_{GT}, \sigma_T, \beta) = N(\mu = Y_{GT}, \sigma^2 = \sigma_T^2 + \beta^2 Y_{GT}^2), \quad (2)$$




Fig. 2. Creating ground-truth using the Dynamic Phantom. During each phantom scan, 200 static volumes were acquired and were averaged voxel-wise to obtain a close approximation to true intensity values. The mean volume was then rotated 600 times synthetically at angles obtained from the optical encoder during the actual run. This yielded ground-truth volumes, which then were compared to the volumes acquired during the scan.




Fig. 3. Qualitative comparison of mean power spectral density of the ground-truth and noise time-series provides signatures for the presence of signal-dependent (non-white) scanner confounds in fMRI data, in addition to the background noise. Voxel-wise noise time-series is calculated by subtracting measured fMRI time-series and the ground-truth time-series.

We estimate the parameters σ_T , and β by Monte-Carlo simulation using Y_{fMRI} and Y_{GT} . Specifically, we model Eq. 1 to sample from the posterior distribution that is proportional to Eq. 2 while assuming constant priors for the parameter distributions. The relative contribution of multiplicative noise to that of the total noise is listed in Table 1 for each scanner. The results indicate that even in modern high-performance scanners with acquisition parameters optimized by a trained MR physicist, the scanner-induced variance due to instability is around 6–18% of the contribution of the total scanner noise. This range is consistent with Greve et al. (2011), in which the authors measured scanner instability by scanning an agar phantom at two varying flip-angles to separate instability from background noise. Because we use different metrics, we included a detailed comparison between the Greve et al. (2011) find-

ings (Supplementary Material, Table 1) and our study in the Supplementary Material. Finally, we provide a case-study comparing the two methods, using modern imaging hardware and acquisition parameters (multi-channel coils and parallel imaging) in the Supplementary Material. We found agreement between the two methods, except when the background noise variance becomes space-variant. This suggests Greve et al.'s method risks inaccuracy for modern acquisition protocols, as previously discussed in Greve et al.

2.4. Using the dynamic phantom generated ground-truth, we quantified scanner-induced non-linearity in fMRI response.

Finally, we observe scanner-induced temporal non-linear distortion of fMRI response using a tree-partition non-linearity estimator (Ljung, 2019) (a piece-wise linear function defined by the binary tree over partitions of the regressor space) with ground-truth as the regressor. Non-linearity is detected in the observed fMRI data if a non-linear function explains significant variance in the observed data beyond the variance explained by the linear function of the ground-truth. Non-linearity estimation was performed using 'isnlrx' function provided in System Identification Toolbox, Matlab (Ljung, 2019), which categorizes non-linearity as strong, weak or not significant based on reliability of the nonlinearity detection test. We observed that the 7T scanner showed the highest non-linearity in response, with 19% of voxels exhibiting strong non-linearity (Table 1).

2.5. Using the dynamic phantom generated ground-truth, we evaluated the efficacy of applying random matrix theory to remove scanner-induced noise; thereby, demonstrating the utility of the dynamic phantom for comparing retrospective denoising techniques against a ground-truth.

A method based on principal component analysis (PCA) coupled with random matrix theory (RMT), called MP-PCA (Veraart et al., 2016a,b),

Table 1

Data-quality metrics for quality control, quantification of scanner-instability, and performance of bandpass filtering, MP-PCA and CNN temporal denoising scheme - evaluated for each scanner using phantom data. CNN denoising increases ST-SNR by ~ 4.7 times the measured time-series, bringing ST-SNR to a regime where signal power is higher than that of noise (ST-SNR 1). Fidelity, ST-SNR and Instability are calculated on a single time series obtained after temporally concatenating voxels of interest.

[illegible]

has been introduced recently for denoising diffusion MRI) (Veraart et al., 2016a,b) and fMRI data (Adhikari et al., 2018). MP-PCA is a 4d image denoising technique that exploits redundancy in the PCA domain using the universal Marchenko–Pastur distribution to remove scanner-induced noise. MP-PCA denoising, followed by bandpass filtering in the frequency-band of interest (0.008-0.1Hz), showed increases in ST-SNR and Dynamic Fidelity over the observed fMRI data and the conventional bandpass filtering (0.008-0.1Hz). MP-PCA denoising showed a significant increase in Dynamic Fidelity with around 40%, 60%, 48%, and 35% increase and a ~2- to 3-fold increase in ST-SNR, for Site 1: PRISMA, Site 2: PRISMA, SKYRA, and MAGNETOM respectively, compared to the observed fMRI data.

2.6. We designed a data-driven temporal filter and observed robust increases in ST-SNR and Dynamic Fidelity of fMRI time-series after denoising.

We provide a deep-learning framework using a Convolutional Neural Network (CNN) for learning an equivalent of a temporal filter. Given that we now have known dynamic inputs, we developed an end-to-end trainable CNN architecture that uses discriminative denoising to remove noise in the hidden layers. We provided pairs of measured fMRI time-series and known signal to learn a mapping from noisy to clean time-series implicitly. We used batch regularization with small batches of batch-size=8 within CNN to avoid internal covariate shift, accelerate the training process, and reduce dependence on network parameter initialization (Sergey Ioffe, 2015). Sigmoid activation function has been used for non-linear mapping and a dropout layer for regularization (Nitish Srivastava, 2014). The architecture details of CNN are specified in Fig. 4.

For evaluating the performance and generalizability of the CNN, we compare the results of CNN denoised fMRI time-series, as shown in Fig. 5, with the original data-quality and temporal de-noising using a standard third-order Butterworth bandpass filter (0.008–0.1 Hz). CNN de-noising showed a significant increase in Dynamic Fidelity with around 53%, 72%, 58%, and 38% increase, for Site 1: PRISMA, Site 2: PRISMA, SKYRA, and MAGNETOM respectively, compared to the observed fMRI data. Further, the CNN de-noising showed a ~ 4 - to 7-fold increase in ST-SNR compared to the observed fMRI data. Finally, CNN de-noising outperforms the conventional temporal bandpass filtering and the MP-PCA denoising (Table 1) in terms of improving both the ST-SNR and the Dynamic Fidelity. While optimal denoising requires collecting both the training and test data during the same session as shown in Table 1, the CNN denoising shows improvement in both ST-SNR and Dynamic Fidelity even with training datasets acquired a few weeks apart from the test dataset (Supplementary Material, Table 4).

2.7. Removing scanner-induced variance from human fMRI data increased the detection sensitivity of brain networks, visible even at the single-subject level.

For assessing the effects of CNN de-noising on human fMRI data, the detection sensitivity of brain networks engaged in movie watching was calculated as a measure of the ability to preserve fluctuations of interest (signal) while removing scanner confounds (noise) from the time-series, and was quantified using the ratio of mean absolute Z-score inside and outside well-defined resting-state network masks in subject-specific ICA maps. A ratio > 1 indicates that Z-score inside the mask is higher compared to voxels outside. Higher this ratio, the easier it is to detect the brain/resting-state networks. We observed an increase in detection sensitivity at the single-subject level for all three scanners after accounting for scanner-related noise, for both the MP-PCA denoising and the CNN denoising method. For MP-PCA denoising, permutation-testing revealed a significant increase in detection sensitivity for all




Fig. 4. Architecture of the convolutional neural network used for discriminative denoising. Each convolution layer (except the last) contains 18 filters with a kernel size of 9 and a stride of 1. Sigmoid is used as the activation function. A dropout of 0.2 is used in the dropout layer. The last convolution layer contains only one filter. Negative of R-squared between the ground-truth and the denoised time-series used as the loss function (minimize) with Adam optimizer for stochastic optimization (Diederik P. Kingma and Ba 2014).




Fig. 5. Exemplar denoising of fMRI output using the trained CNN for two voxels with A. low ST-SNR (0.06), and B. high ST-SNR (0.31) levels.

three scanners (3T PRISMA: percent-change= 9.06% p-value=0.016; 3T SKYRA: percent-change= 13.03%, p-value=0.016; 7T MAGNETOM: percent-change=9.3%, p-value=0.015). Similar trends were observed for CNN denoising (3T PRISMA: percent-change= 13.63% p-value=0.016; 3T SKYRA: percent-change= 20.7%, p-value=0.015; 7T MAGNETOM: percent-change=18.74%, p-value=0.015). Furthermore, the CNN denoising outperformed MP-PCA denoising as evident from Table 2 (3T PRISMA: percent-difference = 4.19% p-value=0.016; 3T SKYRA: percent-difference = 6.78%, p-value=0.03; 7T MAGNETOM: percent-difference =8.64%, p-value=0.015).

3. Discussion

3.1. Why should one use a dynamic phantom rather than a static phantom?

Static phantoms are commonly used for quality assurance (Friedman and Glover, 2006) to assess and minimize scanner fluctuations due to background noise and instability. However, the resting-state fMRI or naturalistic paradigms depend not only upon suppressing fluctuations due to noise but equally upon sensitivity towards signal change, which can only be assessed by a phantom that produces a known and

Table 2

Detection Sensitivity of resting-state networks. Denoising of human fMRI data for removing scanner-confounds showed an increase in the detection sensitivity of resting-state networks. Here, detection sensitivity refers to the ratio of mean absolute Z-score inside and outside a well-defined resting-state network mask in subject-specific ICA maps. Detection sensitivity >1 indicates higher contrast of voxels-of-interest inside the brain network compared to voxels outside. The higher score is in bold.

	Subject #	3T PRISMA			3T SKYRA			7T MAGNETOM		
		Standard Method	MP-PCA Denoising	CNN Denoising	Standard Method	MP-PCA Denoising	CNN Denoising	Standard Method	MP-PCA Denoising	CNN Denoising
Run 1	SUB 1	2.06	2.22	2.24	1.9	2.14	2.27	1.89	2.05	2.25
	SUB 2	2.25	2.45	2.55	2.11	2.39	2.39	2.37	2.61	2.85
	SUB 3	2.17	2.36	2.55	2.13	2.41	2.54	2.32	2.56	2.78
Run 2	SUB 1	1.88	2.01	2.04	1.95	2.25	2.38	2.12	2.24	2.44
	SUB 2	2.23	2.46	2.59	2.32	2.6	2.86	2.22	2.5	2.63
	SUB 3	2.1	2.34	2.45	2.25	2.52	2.84	2.42	2.62	2.89

changing (dynamic) signal. The importance of a dynamic phantom is that it is the only method, to our knowledge, that can quantifiably assess the most basic assumption underlying all task-free fMRI: fidelity between input (brain) dynamics and output (measured fMRI time-series) dynamics. We introduced a novel method for generating ground-truth using the dynamic phantom and estimating voxel-wise noise time-series. The dynamic phantom additionally provides an estimate of standardized signal-to-noise ratio (ST-SNR) and non-linearity, quantifying actual measurement error in fMRI response as compared to static-phantom derived temporal stability of the mean signal (tSNR). While static phantoms estimate only flat-spectrum noise (Expert et al., 2011), the dynamic phantom can detect both signal-dependent and background noise. Using Bayesian parameter estimation, we quantified the ratio of instability/multiplicative noise to the background noise. Although fMRI time-series have several sources of confounds and variance contributed by scanner-instability is relatively small, the reliability of the longitudinal data may be seriously affected without proper characterization. Using data metrics introduced, quality assurance protocols can be established for scanner health monitoring. Any deviations in ST-SNR, Dynamic Fidelity or scanner-instability, compared against longitudinally tracked measurements, would indicate scanner problems.

3.2. Why did we use a deep-learning approach for temporal denoising?

Scanner-instability and background noise in resting-state data lead to decreased detection-sensitivity of resting-state networks, which have been typically addressed by increasing the amount of data collected or increasing the scan-time per subject. These methods are not only expensive but lead to other problems such as subject-fatigue and increased head-motion, which are especially acute in clinical populations. In the current report, we propose a fundamentally different approach for removing scanner confounds from fMRI time-series, which may circumvent the need for collecting more data, and which is ideally suited for single-subject level analyses required for clinical and computational modeling applications, as well as large-scale multi-site and longitudinal studies. Our method exploits the availability of paired measured fMRI and ground-truth data to perform discriminative denoising using CNN. Developing a denoising algorithm for correcting time-series distortions can be framed as a system-identification problem, wherein the goal is to infer a functional relationship between the system input (measured fMRI data) and the system output (denoised fMRI data). Convolution of the measured signal with the identified filter produces the denoised signal. While dealing with linear systems, this system-identification problem reduces to the characterization of impulse response using delta function or observing the system's frequency response using sinusoids. However, for non-linear systems, there exists no canonical representation of the system that will capture "all possibilities" of mapping inputs to transformed outputs. The convolution integral for linear systems can be extended to convolution-like Volterra series for non-linear systems, which can further be extended to Wiener series where each component of the series is orthogonal to all lower-order components. Lee and Schetzen (Lee and Schetzen, 1965) provided a simple method based on cross-correlation for estimating Wiener kernels. However, the cross-correlation method is fundamentally limited by the fact that inputs must be Gaussian. Further, the kernel estimation suffers in cases of strongly nonlinear systems. To overcome these problems, we used deep learning for performing temporal filtering. Intuitively, the trained CNN can be thought of as a temporal filter (like a bandpass-filter), but with filter parameters estimated in an automated data-driven manner optimized for a specific scanner performing a session.

3.3. Why is a dynamic phantom more useful than ICA-based techniques in mitigating scanner-effects for multi-site studies?

Different sites generally have very different-levels of scanner-noise (Greve et al., 2011), causing heteroscedasticity when using ordinary

least-squares estimator and skewing the p-values to be smaller than they should be. Scanner-differences can be reduced by data-processing techniques before analysis (resting-state data), or scanner-effects can be adjusted statistically (task-based data). Feis et al. (Feis et al., 2015) recently showed the successful application of FMRIB's ICA-based X-noiseifier (FIX) (Salimi-Khorshidi et al., 2014) to remove scanner-specific structured noise components that diminished differences in detected resting-state networks across sites. However, the complexity in re-training the FIX classifier for a dataset from every new scanner is non-trivial and requires manual component labeling using data from multiple subjects by an expert. While our measurement of ST-SNR provides a way for statistical adjustment of scanner effects in task-based paradigm using ST-SNR as a covariate in ANCOVA designs, the CNN denoising can remove scanner-induced effects before analysis for resting-state fMRI or naturalistic paradigms in an automated fashion.

3.4. Future directions

Our work has direct implications in moving towards single-subject imaging, which is necessary for clinical purposes as well as for fMRI driven computational neuroscience. Ensuring the stability of time-series adds statistical power to draw useful conclusions from a limited amount of data. Although first-level analysis is dominated by physiological noise (Greve et al., 2011; Triantafyllou et al., 2005; Wald and Polimeni, 2017), we observed a ~13–20% increase in detection sensitivity of resting-state networks after removal of scanner-related noise. The fact that the dynamic phantom can provide, for the first time, a ground truth, permits identification and removal of scanner-related noise. It also enables rigorous evaluation of new data-driven denoising methods under “real-world” conditions that may deviate from idealized *a priori* assumptions (i.e., physical models) of scanner noise characteristics. The dynamic phantom's optical encoder provides precise information (resolution = 0.04392 degrees) about phantom rotation, which can be used for evaluating both prospective and retrospective in-plane motion correction algorithms. Using the dynamic phantom for establishing data-quality metrics, will provide an evaluation of modern imaging protocols, for example compressed sensing fMRI or 3D EPI. Future studies with a larger sample-size will focus on the effects of removing scanner confounds on reliability estimates of functional connectivity analysis and computational neuroscience circuits. Low reliability causes low reproducibility of functional connectomics (Zuo et al., 2019). Reproducibility across sessions while scanning the same patient affects the clinical decision making and thus is an active concern for the use of resting-state fMRI as a clinical tool (O'Connor and Zeffiro, 2019). Further, as physiological noise, thermal noise, and scanner instability are temporally independent, the effect of physiological noise and scanner-induced fluctuations can be regressed out using a general linear model (GLM) framework. The second-order effects/interaction between physiological noise and scanner-induced fluctuations can easily be modeled using interaction terms in the GLM if external physiological recordings are available. The CNN output (denoised fMRI signal) and input (measured fMRI signal) can be used to obtain the regressors (subtracting denoised fMRI signal from the measured fMRI signal) for scanner-induced fluctuations, to model the interaction effects. Additionally, future directions include investigating effects of dynamic phantom estimated ST-SNR on activation effect size in task-based studies, combining multi-site task-based studies using ST-SNR as a covariate, and using CNN denoising to normalize data across sites as required for multi-site studies.

4. Methods

4.1. Study design

We performed imaging at two sites: the SCAN Center at Stony Brook University in Stony Brook, New York (Site 1) and the Athinoula A. Martinos Center for Biomedical Imaging at the Massachusetts General Hospi-

tal in Charlestown, Massachusetts (Site 2). We designed and engineered a dynamic phantom for producing ground-truth time series, based on differences in T_2^* values of agarose gel across voxels of interest. Controlled rotation of the dynamic phantom produces variation in the T_2^* values within a voxel, tuned to generate amplitude changes/signal as observed with BOLD contrast in humans (see Results for a detailed description of the design). At Site 1 (3T Siemens PRISMA scanner), we scanned the phantom during a single session with five acquisition runs, with each successive run separated by a 20-minute interval. Each run had a unique programmed rotation profile as input to the phantom. No human data acquisition occurred at Site 1. At Site 2, we acquired data from three human subjects (two males and one female aged 55, 56, and 47 years, respectively) and the phantom, using three scanners: 3T Siemens SKYRA, 3T Siemens PRISMA, and 7T Siemens MAGNETOM. We acquired data in three imaging sessions: one session per scanner. During each imaging session, we acquired three phantom scans, each with a unique rotation profile, and six human scans, with two scans per subject. The first phantom scan took place at the beginning of each session. Next, each of the three human subjects were scanned while they viewed a naturalistic movie (no audio, see Supplementary Material for video) inside the scanner. Afterward, we acquired the second phantom scan, followed by a repeated acquisition for all three human subjects under identical conditions. Finally, we acquired the third phantom scan. The Institutional Review Board at Massachusetts General Hospital (Partner's Healthcare) provided approval for the human study, and all participants provided written informed consent prior to participating in the study.

4.2. Data acquisition parameters

To ensure that results conservatively reflect actual data-quality metrics within the neuroimaging field, we asked each scanner's MR physicist to independently provide the optimal acquisition parameters for modern fMRI studies conducted on that specific scanner. The details of the protocol parameters are as follows. (1) Site 1 (phantom imaging only): The phantom was scanned on a 3T Siemens PRISMA scanner with a 64-channel head coil. For relaxation rate measurements, multi-echo gradient-echo images were acquired at twelve echo times equally spaced between 5 ms and 60 ms with TR = 70 ms, FOV = 192 mm × 192 mm, flip angle = 20°, slice thickness = 1.5 mm, and readout bandwidth = 320 Hz/px. For the time-series data, standard single-shot gradient-echo EPI data were acquired with the parameters as listed in Table 3. (2) Site 2: Three different scanners were used for data acquisition. For phantom measurement, only EPI scans were acquired. For human measurements, structural scans based on a standard T_1 -weighted MPRAGE and B_0 field maps were acquired in addition to the EPI scans. EPI scan parameters for all three scanners are listed in Table 3. Specifics of structural scans and B_0 field maps are: (a) 3T Siemens SKYRA: Structural scans, for spatial coregistration, were acquired as multi-echo MPRAGE with 1 mm isotropic voxel size and four echoes with TE₁, TE₂, TE₃, TE₄ = 1.69, 3.55, 5.41, 7.27 ms, TR = 2530 ms, flip angle = 7°, and GRAPPA acceleration = 2. B_0 field map images, calculated using phase differences between gradient-echo images at TE = 3.47 ms and 5.93 ms, were acquired (TR = 500 ms, flip angle = 47°, voxel-size = 3.0 × 3.0 × 3.0 mm³ and 44 slices) for EPI distortion correction arising due to susceptibility-induced magnetic field inhomogeneity; (b) 3T Siemens PRISMA: Structural scans were acquired using a single-echo MPRAGE with 1 mm isotropic voxel size, TE = 2.9 ms, TR = 2500 ms, flip angle = 8° and GRAPPA acceleration = 2. B_0 field maps were acquired with TE = 3.47 and 5.93 ms, TR = 500 ms, flip angle = 47°, voxel-size = 3 × 3 × 3 mm and 52 slices; (c) 7T Siemens MAGNETOM: Structural scans were acquired as multi-echo MPRAGE with 1 mm isotropic voxel size at four echoes with TE₁, TE₂, TE₃, TE₄ = 1.61, 3.47, 5.33, 7.19 ms, TR = 2530 ms, flip angle = 7°, and GRAPPA acceleration = 2. B_0 field map images were acquired at TE = 4.60 and 5.62 ms, TR = 723 ms, flip angle = 47°, voxel-size = 1.7 × 1.7 × 1.5 mm³ and 89 slices.

Table 3

Acquisition parameters for functional EPI datasets for both the phantom and human subjects. *Only 600 volumes acquired in the case of human subjects.

Parameter	Site 1		Site 2	
Scanner	Siemens PRISMA	Siemens PRISMA	Siemens SKYRA	Siemens MAGNETOM
B0 Field	3T	3T	3T	7T
Head Coil	64	32	32	32
TR (msec)	1000	800	748	802
TE (msec)	33	30	31	20
Flip Angle (degrees)	52	52	52	33
EPI Factor	84	90	80	96
Voxel Size	2.5mm Isotropic	2.4mm Isotropic	2.5mm Isotropic	2 mm x 2mm x 1.5mm
Number of Slices	28	60	48	85
Number of Volumes*	800/600	800/600	800/600	800/600
Echo-Spacing	0.58	0.51	0.59	0.55
iPAT	1	1	1	2
Multiband Factor	4	6	6	5
Bandwidth (Hz/Px)	2990	2778	2232	2368

4.3. Preprocessing of phantom data for calculating data quality metrics and training the convolutional neural network (CNN)

Acquisition of phantom EPI data involved acquiring the first 200 volumes without any programmed rotation, followed by 600 rotating volumes with the rotation synchronized to the scanner's TR (repetition time) trigger signal. The phantom rotation was limited to around 250 ms from the start of each TR and was quantified using the optical encoder's feedback (Fig. 1A, C). Before analysis, we corrected all phantom acquisitions for smooth spatial intensity variations caused by nonuniformity in the B0 field, B1⁺ field, and receiver coil sensitivity (Sled et al., 1998; Sled and Bruce Pike, 1998) using the N4ITK algorithm (Tustison et al., 2010), implemented in ANTs toolbox. N4ITK offers improved bias field correction over the original nonuniform intensity normalization (N3) algorithm (Sled et al., 1998), via robust b-spline approximation and a hierarchical optimizer to model a range of bias modulation. Additionally, to further minimize the effect of spatial intensity variations for calculating ground truth (Fig. 2), we limit the phantom rotation within 10 degrees clockwise or anti-clockwise relative to the start position at time, $t = 0$. Based on the optical encoder's feedback and scanner's slice timing information, all the slices acquired during in-plane rotation within a TR were discarded from the respective EPI dataset for any further analysis. The remaining slices were manually inspected, and bad slices due to susceptibility artifacts (towards the top and bottom face of the cylinder) were thrown out. The final set of slices then underwent an automated procedure based on contour finding and the Hough transform for generating masks used to select the voxels of interest located in the inner cylinder of each slice. The first 200 volumes of all the remaining slices were averaged voxel-wise to create a mean functional dataset to obtain close approximations to the true voxel intensity. Synthetic rotation of the mean functional dataset, to create ground-truth time-series, involved up-sampling the mean images by a factor of 5 (3rd order spline interpolation), followed by rotation at angles provided by optical encoder's feedback and down-sampling by local averaging to original dimensions of the mean functional slice. Fig. 5 (Supplementary Material) shows that the ground truth signal created using acquisitions at different static positions closely match each other. Subtracting the noisy fMRI output from the corresponding ground-truth time-series yields voxel-wise noise time-series. Power spectrum density (Fig. 3) was calculated using the Welch method implemented in SciPy library (Virtanen et al., 2020). Monte-Carlo simulations for parameter estimation to quantify multiplicative-to-thermal noise ratio were carried out in PyMC3 (Salvatier et al., 2016). We estimated the percentage of voxels exhibiting nonlinearity for each scanner. For a given voxel with ground-truth time series $G(t)$ and a measured fMRI time series $Y(t)$, we express the measured fMRI time series as

$$Y(t) = L(t) + F(t) + E(t)$$

where $L(t)$ represents the portion of data explained by a linear function of the ground-truth time series, $F(t)$ represents the portion of data explained by a nonlinear function of the ground-truth time series and $E(t)$ represents unexplained residual variance. If the nonlinear function explains a significant portion of variance after regressing out the linear model $L(t)$ from $Y(t)$, a nonlinearity is detected in the time series $Y(t)$ (Ljung et al., 2006; Sjöberg et al., 1995). $F(t)$ models the nonlinearity based on a nonlinear function/estimator (Sjöberg et al., 1995; Ljung et al., 2006) of the ground-truth time series, which can be a binary partition tree, a radial basis function network based on wavelets, a piecewise linear estimator, a multi-layer neural network or custom-built non-linearity regressors (for example, quadratic or polynomial regressors of ground-truth time series). We used a binary tree partition (Vanli and Kozat, 2014; Ljung et al., 2006) as nonlinearity estimator, which splits the data into two subsets followed by iterative splitting of each subset into smaller subsets to partition the entire regressor space (ground-truth time series) into a binary tree. After this, linear regression is performed at each level of the binary tree to complete the estimation procedure (Vanli and Kozat, 2014). We performed the nonlinearity estimation with a binary partition tree using the "isnlrx" function provided in System Identification Toolbox, Matlab (Ljung, 2019).

For all voxels, the measured and the ground-truth time-series pairs were used for end-to-end training of the CNN (see Fig. 4 for architecture). Given that multiple phantom scans were acquired for each scanner, CNN training involved combining data acquired with different programmed motion sequences (Supplementary Materials: Suppl. Fig. 2) on a scanner for data-augmentation. Within each training dataset, 33% of data was used as validation split and model weights with lowest validation loss was saved as the trained CNN. For Site 1, three CNNs were trained using data from scans 1 and 3, scans 2 and 4, and scans 3 and 5. For Site 2, three phantom scans were acquired at each scanner, and CNNs were trained using data from scans 1 and 2, scans 2 and 3, and scans 1 and 3. For testing denoising performance, the test data were denoised using a trained CNN which did not use the test data during training (out-of-sample denoising), for example: at Site 2, for denoising scan 2, we used a CNN trained on scans 1 and 3.

4.4. Preprocessing of human data

Spatial preprocessing was performed in the Statistical Parametric Mapping (SPM12) software package (<http://www.fil.ion.ucl.ac.uk/spm>) using the pipeline provided in the CONN toolbox (Whitfield-Gabrieli and Nieto-Castanon, 2012). Functional images were motion (rigid alignment, six-degrees-of-freedom) and B₀ field map corrected, and a mean functional image was calculated for each subject. The mean functional images were then co-registered to high-resolution structural images followed by segmentation to generate gray matter, white matter, and cerebrospinal fluid images. Each voxel time-series was de-

measured and underwent quadratic de-trending. For further temporal preprocessing, the data went through two different pipelines to generate three datasets as discussed below: (a) **Standard Method**: Physiological confounds were removed using the Component-Based Noise Correction Method (Behzadi et al., 2007) (CompCor) implemented through Nipype interface (Gorgolewski et al., 2011). CompCor regresses out the confounding effects of multiple empirically estimated noise sources calculated from variability in BOLD time-series of cerebrospinal fluid and white matter (based on principal component analysis). Five components of white matter and cerebrospinal fluid, and six motion parameters, along with temporal bandpass filtering (0.008–0.1 Hz), were used for physiological denoising. Removal of confounds was orthogonal to the bandpass filtering (Lindquist et al., 2019); (b) **CNN Denoising**: Spatially preprocessed functional data (motion and fieldmap corrected and normalized to MNI) underwent denoising (voxels in gray-matter only) using trained scanner-specific CNN, followed by physiological confound removal as in the standard method (CompCor, motion, and bandpass filtering); (c) **MP-PCA denoising**: We repeated the spatial preprocessing of functional data and applied the standard method of temporal preprocessing, on MP-PCA denoised raw functional data, to generate a third dataset in addition to the standard method and CNN denoising datasets. Finally, datasets obtained from all three denoising methods were smoothed with a 4-mm full width at half-maximum Gaussian kernel, followed by normalization to $2 \times 2 \times 2$ mm Montreal Neurological Institute (MNI) EPI template.

4.5. Calculating detection sensitivity of resting-state networks

To identify functionally connected networks in a data-driven manner, we performed group spatial ICA on the preprocessed data using the GIFT v3.0b fMRI Toolbox (<https://trendscenter.org/software/>), separately for each scanner and temporal processing scheme (standard method and CNN denoising). For each dataset, 20 independent components were obtained, after ten runs of ICASSO (Himberg et al., 2004) procedure for ensuring component stability. Subject-specific spatial maps and associated time courses were estimated using back-reconstruction (GICA) (Erhardt et al., 2011). We used the Infomax algorithm for performing ICA. ICA spatial maps were converted to Z values. We spatially matched the subject-specific ICA maps to seventeen well-defined resting-state network templates obtained from Yeo et al. (Yeo et al., 2011), for obtaining each subject's corresponding network ICA maps. Detection sensitivity was then calculated as the ratio of mean absolute Z-score inside and outside of each of the seventeen resting-state network masks applied to the matched subject-specific ICA spatial maps. The mean of detection sensitivity values, across all seventeen networks, for each subject, yielded a total of six values (three subjects with two runs) for every scanner. These six values were compared between the standard method, MP-PCA denoising, and the CNN temporal denoising for each scanner using permutation testing (100,000 repetitions).

Declaration of Competing Interest

R.K, A.L, J.R.P declare no conflict of interest. The intellectual property covering the dynamic phantom is owned by the Research Foundation for The State University of New York (RF SUNY). L.R.M.-P and H.H.S, as inventors, will receive a share of the royalties under the terms prescribed by the RF SUNY. ALA Scientific Inc., NY – USA, is the licensee for commercialization of the dynamic phantom. L.T declares employment, and A.K declares employment & ownership in ALA Scientific Inc.

Credit authorship contribution statement

Rajat Kumar: Conceptualization, Methodology, Software, Formal analysis, Investigation, Resources, Writing - original draft, Writing - review & editing. **Liang Tan**: Software, Resources. **Alan Kriegstein**: Resources, Funding acquisition. **Andrew Lithen**: Investiga-

tion. **Jonathan R. Polimeni**: Supervision, Writing - review & editing. **Lilianne R. Mujica-Parodi**: Conceptualization, Methodology, Supervision, Writing - original draft, Writing - review & editing, Project administration, Funding acquisition. **Helmut H. Strey**: Conceptualization, Methodology, Software, Resources, Supervision, Writing - original draft, Writing - review & editing, Funding acquisition.

Acknowledgements

The research described in this paper was funded by the National Science Foundation (STTR Phase 1 Grant 1622525), the National Institute of Drug Abuse (SBIR Phase 1 & Phase 2 Grant 1R44 DA043277-01) and the National Heart, Lung and Blood Institute (Grant 5U01HL12752202). The authors gratefully acknowledge the assistance provided by Jie Ma and Michael Bishop during design and engineering of the dynamic phantom, and Anar Amgalan for discussions on methodology and software development.

Data and code availability statement

The data used in this study are available from OpenNeuro upon request to the corresponding authors, without requiring any formal data-sharing agreement. For data analysis, SPM can be obtained from - <https://www.fil.ion.ucl.ac.uk/spm/>; Conn-Toolbox can be obtained from - <https://web.conn-toolbox.org/>; GIFT can be obtained from - <https://trendscenter.org/software/>; custom-software code used in the analysis can be obtained from - https://github.com/RajatKGupta/fMRI_BrainDancer.

Supplementary Material

Supplementary Material associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2020.117584.

References

- Adhikari, Bhim M., Jahanshad, Neda, Shukla, Dinesh, Glahn, David C., Blangero, John, Reynolds, Richard C., Cox, Robert W., Fieremans, Els, Veraart, Jelle, Novikov, Dmitry S., Nichols, Thomas E., Elliot Hong, L., Thompson, Paul M., Kochunov, Peter, 2018. 'Heritability estimates on resting state fMRI data using ENIGMA analysis pipeline'. *Pac. Symp. Biocomput.* 23, 307–318.
- Behzadi, Yashar, Restom, Khaled, Liu, Joy, Liu, Thomas T., 2007. 'A component based noise correction method (CompCor) for BOLD and perfusion based fMRI'. *NeuroImage* 37, 90–101.
- Bollmann, Saskia, Lars Kasper, S., Vannesjo, Johanna, Diaconescu, Andreea O., Dietrich, Benjamin E., Gross, Simon, Stephan, Klaas E., Pruessmann, Klaas P., 2017. Analysis and correction of field fluctuations in fMRI data using field monitoring. *NeuroImage* 154, 92–105.
- Caballero-Gaudes, César, Reynolds, Richard C., 2017. 'Methods for cleaning the BOLD fMRI signal'. *NeuroImage* 154, 128–149.
- Ciuciu, Philippe, Varoquaux, Gaël, Abry, Patrice, Sadaghiani, Sepideh, Kleinschmidt, Andreas, 2012. 'Scale-free and multifractal properties of fMRI signals during rest and task'. *Front. Physiol.* 3.
- DeDora, Daniel J., Nedic, Sanja, Katti, Pratha, Arnab, Shafique, Wald, Lawrence L., Takahashi, Atsushi, Dijk, Koene R.A., Van, Strey, Helmut H., Mujica-Parodi, Lilianne R., 2016. 'Signal Fluctuation Sensitivity: An Improved Metric for Optimizing Detection of Resting-State fMRI Networks'. *Front. Neurosci.* 10, 180–80.
- Diederik P. Kingma, and Jimmy Ba. 2014. 'Adam: A Method for Stochastic Optimization', arXiv:1412.6980.
- Djabourov, M., Clark, A.H., Rowlands, D.W., Ross-Murphy, S.B., 1989. 'Small-angle x-ray scattering characterization of agarose gels and gels'. *Macromolecules* 22, 180–188.
- Duerst, Yolanda, Wilm, Bertram J., Dietrich, Benjamin E., Johanna Vannesjo, S., Barmet, Christoph, Schmid, Thomas, Brunner, David O., Pruessmann, Klaas P., 2015. 'Real-time feedback for spatiotemporal field stabilization in MR systems'. *Magn. Reson. Med.* 73, 884–893.
- Eklund, Anders, Thomas E. Nichols, and Hans Knutsson. 2016. 'Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates', 113: 7900–05.
- Erhardt, Erik Barry, Rachakonda, Srinivas, Bedrick, Edward J., Allen, Elena A., Adali, Tülay, Calhoun, Vince D., 2011. Comparison of multi-subject ICA methods for analysis of fMRI data. *Hum. Brain Map.* 32, 2075–2095.
- Expert, Paul, Lambiotte, Renaud, Chialvo, Dante R., Christensen, Kim, Jeldtoft Jensen, Henrik, Sharp, David J., Turkheimer, Federico, 2011. 'Self-similar correlation function in brain resting-state functional magnetic resonance imaging'. *J. R. Soc. Interface* 8, 472–479.

- Feis, Rogier A., Smith, Stephen M., Filippini, Nicola, Douaud, Gwenaëlle, Dopfer, Elise G.P., Heise, Verena, Trachtenberg, Aaron J., van Swieten, John C., van Buchem, Mark A., Rombouts, Serge A.R.B., Mackay, Clare E., 2015. ICA-based artifact removal diminishes scan site differences in multi-center resting-state fMRI. *Front. Neurosci.* 9.
- Friedman, Lee, Glover, Gary H., 2006. 'Report on a multicenter fMRI quality assurance protocol'. *J. Magn. Reson. imaging: JMIR* 23, 827–839.
- Friedman, Lee, Stern, Hal, Brown, Gregory G., Mathalon, Daniel H., Turner, Jessica, Glover, Gary H., Gollub, Randy L., Lauriello, John, Lim, Kelvin O., Cannon, Tyrone, Greve, Douglas N., Bockholt, Henry Jeremy, Belger, Aysenil, Mueller, Bryon, Doty, Michael J., He, Jianchun, Wells, William, Smyth, Padhraic, Pieper, Steve, Kim, Seyoung, Kubicki, Marek, Vangel, Mark, Potkin, Steven G., 2008. Test–retest and between-site reliability in a multicenter fMRI study. *Hum. Brain Mapp.* 29, 958–972.
- Glover, Gary H., Mueller, Bryon A., Turner, Jessica A., van Erp, Theo G.M., Liu, Thomas T., Greve, Douglas N., Voyvodic, James T., Rasmussen, Jerod, Brown, Gregory G., Keator, David B., Calhoun, Vince D., Lee, Hyo Jong, Ford, Judith M., Mathalon, Daniel H., Diaz, Michele, O'Leary, Daniel S., Gadde, Syam, Preda, Adrian, Lim, Kelvin O., Wible, Cynthia G., Stern, Hal S., Belger, Aysenil, McCarthy, Gregory, Ozyurt, Burak, Potkin, Steven G., 2012. 'Function biomedical informatics research network recommendations for prospective multicenter functional MRI studies'. *J. Magn. Reson. Imaging: JMIR* 36, 39–54.
- Gorgolewski, Krzysztof, Burns, Christopher, Madison, Cindee, Clark, Dav, Halchenko, Yaroslav, Waskom, Michael, Ghosh, Satrajit, 2011. 'Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python'. *Front. Neuroinform.* 5.
- Greve, Douglas N., Brown, Gregory G., Mueller, Bryon A., Glover, Gary, Liu, Thomas T., Psychometrika, Function Biomedical Research Network %J, 2013. 'A Survey of the Sources of Noise in fMRI'. *Psychometrika* 78, 396–416.
- Greve, Douglas N., Mueller, Bryon A., Liu, Thomas, Turner, Jessica A., Voyvodic, James, Yetter, Elizabeth, Diaz, Michele, McCarthy, Gregory, Wallace, Stuart, Roach, Brian J., Ford, Judy M., Mathalon, Daniel H., Calhoun, Vince D., Wible, Cynthia G., Brown, Gregory G., Potkin, Steven G., Glover, Gary, 2011. 'A novel method for quantifying scanner instability in fMRI'. *Magn. Reson. Med.* 65, 1053–1061.
- Himberg, Johan, Hyvärinen, Aapo, Esposito, Fabrizio, 2004. 'Validating the independent components of neuroimaging time series via clustering and visualization'. *NeuroImage* 22, 1214–1222.
- Jezzard, Peter, Clare, Stuart, 1999. 'Sources of distortion in functional MRI data'. *Hum. Brain Mapp.* 8, 80–85.
- Kasper, Lars, Bollmann, Saskia, Vannesjo, S.Johanna, Gross, Simon, Haeblerlin, Maximilian, Dietrich, Benjamin E., Pruessmann, Klaas P., 2015. Monitoring, analysis, and correction of magnetic field fluctuations in echo planar imaging time series. *Magn. Reson. Med.* 74, 396–409.
- Lauterbur, Zhi-Pei Liang, Paul, C., 2000. *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. Wiley-IEEE Press.
- Lee, M.H., Smyser, C.D., Shimony, J.S., 2013. 'Resting-state fMRI: a review of methods and clinical applications'. *Am. J. Neuroradiol.* 34, 1866–1872.
- Lee, Y.W., Schetzen, M., 1965. 'Measurement of the Wiener Kernels of a non-linear system by cross-correlation'. *Int. J. Control* 2, 237–254.
- Lindquist, Martin A., Geuter, Stephan, Wager, Tor D., Caffo, Brian S., 2019. 'Modular preprocessing pipelines can reintroduce artifacts into fMRI data'. *Hum. Brain Mapp.* 40, 2358–2376.
- Liu, Thomas T., 2016. 'Noise contributions to the fMRI signal: An overview'. *NeuroImage* 143, 141–151.
- Ljung, Lennart, 2019. *System Identification Toolbox Reference*. The MathWorks, Inc..
- Ljung, Lennart, Zhang, Qinghua, Lindskog, Peter, Iouditski, Anatoli, Singh, Rajiv, 2006. 'An integrated system identification toolbox for linear and non-linear models'. *IFAC Proc. Vol.* 39, 931–936.
- Mujica-Parodi, Lilianne R., Cha, Jiook, Gao, Jonathan, 2017. 'From anxious to reckless: a control systems approach unifies prefrontal-limbic regulation across the spectrum of threat detection'. *Front. Syst. Neurosci.* 11.
- Srivastava, Nitish, Hinton, Geoffrey, Krizhevsky, Alex, Sutskever, Ilya, Salakhutdinov, Ruslan, 2014. 'Dropout: A Simple Way to Prevent Neural Networks from Overfitting'. *J. Mach. Learn. Res.* 15, 1929–1958.
- O'Connor, Erin E., Zeffiro, Thomas A., 2019. 'Why is clinical fMRI in a resting state?'. *Front. Neurol.* 10.
- Power, Jonathan D., Barnes, Kelly A., Snyder, Abraham Z., Schlaggar, Bradley L., Petersen, Steven E., 2012. 'Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion'. *NeuroImage* 59, 2142–2154.
- Rădulescu, Anca, Mujica-Parodi, Lilianne R., 2014. 'Network connectivity modulates power spectrum scale invariance'. *NeuroImage* 90, 436–448.
- Salimi-Khorshidi, Gholamreza, Douaud, Gwenaëlle, Beckmann, Christian F., Glasser, Matthew F., Griffanti, Ludovica, Smith, Stephen M., 2014. Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. *NeuroImage* 90, 449–468.
- Salvatier, John, Wiecki, Thomas V., Fonnesbeck, Christopher, 2016. 'Probabilistic programming in Python using PyMC3'. *PeerJ Comput. Sci.*
- Sergey Ioffe, Christian Szegedy, 2015. 'Batch normalization: accelerating deep network training by reducing internal covariate shift', arXiv:1502.03167.
- Sjöberg, Jonas, Zhang, Qinghua, Ljung, Lennart, Benveniste, Albert, Delyon, Bernard, Glorennec, Pierre-Yves, Hjalmarsson, Håkan, Juditsky, Anatoli, 1995. 'Nonlinear black-box modeling in system identification: a unified overview'. *Automatica* 31, 1691–1724.
- Sled, J.G., Zijdenbos, A.P., Evans, A.C., 1998. 'A nonparametric method for automatic correction of intensity nonuniformity in MRI data'. *IEEE Trans. Med. imaging* 17, 87–97.
- edited by Sled, John G., Bruce Pike, G., 1998. Understanding intensity non-uniformity in MRI. In: Wells, William M., Colchester, Alan, Delp, Scott (Eds.), *Medical Image Computing and Computer-Assisted Intervention — MICCAI'98*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 614–622 edited by.
- Triantafyllou, C., Hoge, R.D., Krueger, G., Wiggins, C.J., Potthast, A., Wiggins, G.C., Wald, J.L., 2005. Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *NeuroImage* 26, 243–250.
- Tustison, Nicholas J., Avants, Brian B., Cook, Philip A., Zheng, Yuanjie, Egan, Alexander, Yushkevich, Paul A., Gee, James C., 2010. N4ITK: improved N3 bias correction'. *IEEE Trans. Med. Imaging* 29, 1310–1320.
- Van Horn, John Darrell, Toga, Arthur W., 2009. 'Multisite neuroimaging trials'. *Curr. Opin. Neurol.* 22, 370–378.
- Vanli, N.D., Kozat, S.S., 2014. 'A Comprehensive approach to universal piecewise nonlinear regression based on trees'. *IEEE Trans. Signal Process.* 62, 5471–5486.
- Veraart, Jelle, Fieremans, Els, Novikov, Dmitry S., 2016. 'Diffusion MRI noise mapping using random matrix theory'. *Magn. Reson. Med.* 76, 1582–1593.
- Veraart, Jelle, Novikov, Dmitry S., Christiaens, Daan, Ades-Aron, Benjamin, Sijbers, Jan, Fieremans, Els, 2016. 'Denoising of diffusion MRI using random matrix theory'. *NeuroImage* 142, 394–406.
- Virtanen, Pauli, Gommers, Ralf, Oliphant, Travis E., Haberland, Matt, Reddy, Tyler, Cournapeau, David, Burovski, Evgeni, Peterson, Pearu, Weckesser, Warren, Bright, Jonathan, van der Walt, Stéfan J., Brett, Matthew, Wilson, Joshua, Jarrod Millman, K., Mayorov, Nikolay, Nelson, Andrew R.J., Jones, Eric, Kern, Robert, Larson, Eric, Carey, C.J., Polat, İlhan, Feng, Yu, Moore, Eric W., VanderPlas, Jake, Laxalde, Denis, Perktold, Josef, Cimrman, Robert, Henriksen, Ian, Quintero, E.A., Harris, Charles R., Archibald, Anne M., Ribeiro, Antônio H., Pedregosa, Fabian, van Mulbregt, Paul, Vijaykumar, Aditya, Bardelli, Alessandro Pietro, Rothberg, Alex, Hilboll, Andreas, Kloeckner, Andreas, Scopatz, Anthony, Lee, Antony, Rokem, Ariel, Woods, C.Nathan, Fulton, Chad, Masson, Charles, Hägerström, Christian, Fitzgerald, Clark, Nicholson, David A., Hagen, David R., Pasechnik, Dmitrii V., Olivetti, Emanuele, Martin, Eric, Wieser, Eric, Silva, Fabrice, Lenders, Felix, Wilhelm, Florian, Young, G., Price, Gavin A., Ingold, Gert-Ludwig, Allen, Gregory E., Lee, Gregory R., Audren, Hervé, Probst, Irvin, Dietrich, Jörg P., Silterra, Jacob, Webber, James T., Slavič, Janko, Nothman, Joel, Buchner, Johannes, Kulick, Johannes, Schönberger, Johannes L., Cardoso, José Vinícius de Miranda, Reimer, Joscha, Harrington, Joseph, Rodríguez, Juan Luis Cano, Nunez-Iglesias, Juan, Kuczynski, Justin, Tritz, Kevin, Thoma, Martin, Newville, Matthew, Kümmerer, Matthias, Bolingbroke, Maximilian, Tartre, Michael, Pak, Mikhail, Smith, Nathaniel J., Nowaczyk, Nikolai, Shebanov, Nikolay, Pavlyk, Oleksandr, Brodtkorb, Per A., Lee, Perry, McGibbon, Robert T., Feldbauer, Roman, Lewis, Sam, Tygier, Sam, Sievert, Scott, Vigna, Sebastiano, Peterson, Stefan, More, Surhud, Pudlik, Tadeusz, Oshima, Takuya, Pingel, Thomas J., Robitaille, Thomas P., Spura, Thomas, Jones, Thouis R., Cera, Tim, Leslie, Tim, Zito, Tiziano, Krauss, Tom, Upadhyay, Utkarsh, Halchenko, Yaroslav O., Vázquez-Baeza, YoshikiContributors SciPy, 2020. *SciPy 1.0: fundamental algorithms for scientific computing in Python*. *Nat. Methods* 17, 261–272.
- Wald, Lawrence L., Polimeni, Jonathan R., 2017. 'Impacting the effect of fMRI noise through hardware and acquisition choices – Implications for controlling false positive rates'. *NeuroImage* 154, 15–22.
- Welvaert, Marijke, Rosseel, Yves, 2013. 'On the definition of signal-to-noise ratio and contrast-to-noise ratio for fMRI data'. *PLoS one* 8, e77089.
- Whitfield-Gabrieli, Susan, Nieto-Castanon, Alfonso, 2012. 'Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks'. *Brain Connect.* 2, 125–141.
- Yeo, B.T.Thomas, Krienen, Fenna M., Sepulcre, Jorge, Sabuncu, Mert R., Lashkari, Dania, Hollinshead, Marisa, Roffman, Joshua L., Smoller, Jordan W., Zöllei, Lilla, Polimeni, Jonathan R., Fischl, Bruce, Liu, Hesheng, Buckner, Randy L., 2011. 'The organization of the human cerebral cortex estimated by intrinsic functional connectivity'. *J. Neurophysiol.* 106, 1125–1165.
- Zaitsev, Maxim, Akin, Burak, LeVan, Pierre, Knowles, Benjamin R., 2017. 'Prospective motion correction in functional MRI'. *NeuroImage* 154, 33–42.
- Zuo, Xi-Nian, Biswal, Bharat B., Poldrack, Russell A., 2019. 'Editorial: Reliability and Reproducibility in Functional Connectomics'. *Front. Neurosci.* 13.