

Macroscopic resting-state brain dynamics are best described by linear models

Erfan Nozari^{1,2,3}, Maxwell A. Bertolero⁴, Jennifer Stiso^{4,5}, Lorenzo Caciagli⁴,
Eli J. Cornblath^{4,5}, Xiaosong He⁴, Arun S. Mahadevan⁴, George J. Pappas⁶, and
Dani S. Bassett^{4,6,7,8,9,10,*}

¹*Department of Mechanical Engineering, University of California, Riverside, CA 92521, USA*

²*Department of Electrical and Computer Engineering, University of California, Riverside, CA 92521, USA*

³*Department of Bioengineering, University of California, Riverside, CA 92521, USA*

⁴*Department of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104, USA*

⁵*Department of Neuroscience, University of Pennsylvania, Philadelphia, PA 19104, USA*

⁶*Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA*

⁷*Department of Physics & Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA*

⁸*Department of Neurology, University of Pennsylvania, Philadelphia, PA 19104, USA*

⁹*Department of Psychiatry, University of Pennsylvania, Philadelphia, PA 19104, USA*

¹⁰*Santa Fe Institute, Santa Fe, NM 87501, USA*

** Corresponding author, dsb@seas.upenn.edu*

Abstract

It is typically assumed that large networks of neurons exhibit a large repertoire of nonlinear behaviours. Here we challenge this assumption by leveraging mathematical models derived from measurements of local field potentials via intracranial electroencephalography and of whole-brain blood-oxygen-level-dependent brain activity via functional magnetic resonance imaging. We used state-of-the-art linear and nonlinear families of models to describe spontaneous resting-state activity of 700 participants in the Human Connectome Project and 122 participants in the Restoring Active Memory project. We found that linear autoregressive models provide the best fit across both data types and three performance metrics: predictive power, computational complexity and the extent of the residual dynamics unexplained by the model. To explain this observation, we show that microscopic nonlinear dynamics can be counteracted or masked by four factors associated with macroscopic dynamics: averaging over space and over time, which are inherent to aggregated macroscopic brain activity, and observation noise and limited data samples, which stem from technological limitations. We therefore argue that easier-to-interpret linear models can faithfully describe macroscopic brain dynamics during resting-state conditions.

One-sentence editorial summary (to appear right below the title of your Article on the journal’s website):
Linear mathematical models derived from measurements of local field potentials and of whole-brain blood-oxygen-level-dependent neural activity predict resting-state neural dynamics at least as accurately as nonlinear models.

Throughout the recent history of neuroscience, computational models have been developed and used ubiquitously in order to decompose the complex neural mechanisms underlying cognition and behavior.^{1–5} A dilemma that is inherent to computational modeling, but particularly challenging in computa-

tional neuroscience, is the trade-off between (cross-validated) accuracy and simplicity. Both finely detailed models⁶ and broadly simplified ones^{7,8} have their respective proponents. One of the many facets of this trade-off pertains to the use of linear *versus* nonlinear models. Nonlinearity of dynamics is in-

evitable at the (micro) scale of individual neurons⁹ and their components,¹⁰ and has been demonstrated, though less comprehensively, at the (meso) scale of neuronal populations.¹¹ Further supported by theoretical derivations¹² and motivated by the much larger repertoire of behaviors of nonlinear systems (including chaos, multi-stability, and meta-stability), an assumption has thus formed^{13–16} that accurate models of neurodynamics at the macroscale (i.e., of brain regions) must inevitably be nonlinear.

This assumption begs the question of whether nonlinear models will in fact perform better than linear ones in accounting for the dynamics of neuroimaging and neurophysiological data. Specifically, can nonlinear models explain neuroimaging or neurophysiological data more accurately than linear ones? This pragmatic modeling question, importantly, is different from the general question of whether any *signs* of “nonlinearity” can be found in neuroimaging time series^{17–19} (see Discussion).

Few investigations^{20–22} have indeed sought to answer the former question directly by comparing the “fit” of linear and nonlinear models to neurophysiological (electroencephalography, EEG, and intracranial electroencephalography, iEEG) time series. However, even these few works are limited in that each provides a single comparison between a linear and a nonlinear family of models (linear autoregressive (AR) vs. nonlinear manifold-based model in,²⁰ linear finite impulse response vs. nonlinear Volterra series model in,²¹ and linear state space vs. nonlinear AR with radial basis function nonlinearities in²²), which need not be the best representatives of linear and nonlinear models in general. While the compared linear and nonlinear models were found to be as predictive of EEG data in²⁰ and iEEG data in,²² the results were mixed in.²¹ Using scalp EEG data from patients with epilepsy, this paper finds mostly linear dynamics well within and outside of seizures, and mostly nonlinear dynamics (though varying across patients) in the periods around seizure onsets and offsets. While limited in their scope, these works beg for a deep and rigorous data-driven investigation into the nonlinearity of macroscopic brain dynamics, as pursued herein.

In the second part of the paper, we seek to answer the question of why nonlinear models do not provide more accurate predictions than linear ones despite the fact that neurodynamics are inevitably nonlinear at the microscale. Specifically, we numerically

demonstrate, using a simple sigmoidal nonlinearity, that four properties of macroscopic brain dynamics can fundamentally *counteract* or seemingly *mask* nonlinear dynamics present at the microscale: averaging over the activity of large populations of neurons to obtain a single macroscopic time series (averaging over space), natural low pass-filtering properties of brain processes (averaging over time), observation noise, and limited data samples. While the effects of observation noise and limited data samples are technology dependent but otherwise independent from the form of nonlinearity, the effects of spatiotemporal averaging are fundamental to macroscopic neural dynamics, and may depend on the functional form of the microscale nonlinearity. We thus also verify and demonstrate the effects of spatiotemporal averaging using a data-driven and biophysically grounded spiking neuron model.²³ Together, our results provide important evidence that linear models can be as descriptive as nonlinear ones at the macroscale, as well as a methodology based on system identification theory to quantitatively define a “best” model of whole brain dynamics given a set of specified costs.

Results

System identification and data-driven computational modeling. Among the diverse categories of computational models used in neuroscience, we focus on ordinary differential equation (ODE) models of the general form

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t)) + \mathbf{e}_1(t), \quad \mathbf{x}(0) = \mathbf{x}_0 \quad (1a)$$

$$\mathbf{y}(t) = h(\mathbf{x}(t)) + \mathbf{e}_2(t) \quad (1b)$$

where $\mathbf{y}(t)$ is an n -dimensional time series of recorded brain activity, in this case via resting state fMRI (rsfMRI) or iEEG (rsiEEG), $\mathbf{x}(t)$ is an m -dimensional time series of “internal” brain states, f and h are generally nonlinear vector fields, and $\mathbf{e}_1(t)$ and $\mathbf{e}_2(t)$ are time series of process and measurement noise with arbitrary statistics (**Fig. 1a-b**). It is possible, although not necessary, that $m = n$. As with any differential equation, the description would not be complete without the initial condition $\mathbf{x}(0) = \mathbf{x}_0$, determining the state of the brain at the initial recording time $t = 0$. Note that no external input $\mathbf{u}(t)$ is considered due to the resting state condition of the experiments. Also, given that we can only sample $\mathbf{y}(t)$ in discrete

time, we implement equation (1) by approximating the derivative $\dot{\mathbf{x}}(t)$ as a first difference $\mathbf{x}(t) - \mathbf{x}(t-1)$ (see Methods).

The model in equation (1) is “linear” if the functions f and h are linear functions, i.e., matrix operations of the form $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$ and $h(\mathbf{x}) = \mathbf{C}\mathbf{x}$ where $\mathbf{A}_{m \times m}$ and $\mathbf{C}_{n \times m}$ are constant (or even time-varying, but state-independent) matrices. Throughout the field of computational neuroscience, numerous models of the form in equation (1) or its discretization (see Methods) are constructed and used, each with different functional forms and noise statistics.^{6, 7, 9, 12, 15, 22, 24} The critical but fairly overlooked quest of system identification²⁵ is then to find the “best” model, among all the available options, against experimental data. This comparison, indeed, depends on one’s measure of a model’s goodness of fit.

A natural choice, referred to as the prediction error (PE) approach, is based on how well a given model can predict the *future* of the measured time series from its *current and past* values (**Fig. 1c**). Note that this prediction is precisely what an ODE such as equation (1a) defines: it models the *change* $\dot{\mathbf{x}}(t)$, and thus the immediate future, of the system’s state (and therefore output) from its current state $\mathbf{x}(t)$. Since the state $\mathbf{x}(t)$ is not directly available, it should in turn be estimated from the current and past measurements of the output $\mathbf{y}(t)$. Therefore, the PE approach in its simplest form seeks to minimize, within any given parametric or non-parametric family of models, the magnitude of the one-step ahead PE

$$\varepsilon(t) = \mathbf{y}(t) - \hat{\mathbf{y}}(t|t-1) \quad (2)$$

where $\hat{\mathbf{y}}(t|t-1)$ is the Bayes-optimal, minimum variance estimate of $\mathbf{y}(t)$ given all of the history $\{\mathbf{y}(0), \dots, \mathbf{y}(t-1)\}$ of \mathbf{y} up to time $t-1$ (Ref.²⁵) (**Fig. 1c**). Notably, the PE approach focuses on the prediction accuracy of the time series itself, rather than the prediction accuracy of functional connectivity (**Fig. 1d**) or other statistics of the time series (cf. Discussion and **Supplementary Fig. 1**). This approach can also be easily extended to multi-step ahead predictions, cf.²⁵ and **Supplementary Figs. 20,21**.

The task of system identification does not end once the parameters of a given family of models are fit to the (training) data. The critical next step is to assess the quality of the fit, particularly to data withheld

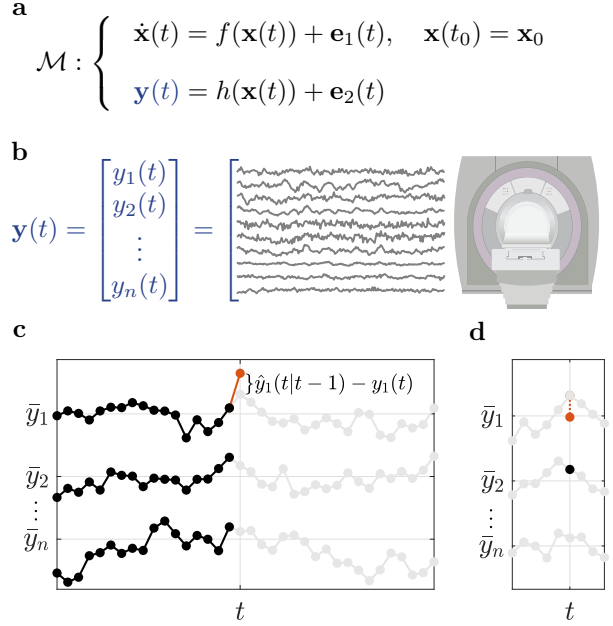


Fig. 1 | Prediction error method for system identification. (a) The general category of computational models \mathcal{M} studied in this work, represented by an ODE describing the resting state evolution of internal states $\mathbf{x}(t)$ and an output model that maps internal states to fMRI/iEEG time series $\mathbf{y}(t)$, as shown for fMRI in panel (b). A total of $n_{\text{fMRI}} = 116$ regions are used throughout (see Methods) for fMRI, while $13 \leq n_{\text{iEEG}} \leq 175$ channels are used for each iEEG patient. (c) A schematic representation of the prediction error system identification framework used in this work. At each time t , all of $\{\mathbf{y}(0), \dots, \mathbf{y}(t-1)\}$ is used to predict $\mathbf{y}(t)$, simultaneously across all channels, which is denoted for simplicity by $\hat{\mathbf{y}}(t|t-1)$. \bar{y}_i denotes the temporal average of $y_i(t)$ for each channel i . This estimate should not be confused with estimates of functional connectivity (FC). (d) FC measures the covariation between pairs of channels or, equivalently, how well each $y_i(t)$ ($y_1(t)$ in the figure) can be predicted from each other $y_j(t)$ ($y_2(t)$ in the figure), at the same time t .

during the training (cross-validation). In the PE approach, the two most widely used measures are the variance and the whiteness of the PE (Ref.²⁵), where the former is often measured by

$$R_i^2 = 1 - \frac{\sum_t \varepsilon_i(t)^2}{\sum_t (y_i(t) - \bar{y}_i)^2} \quad (3)$$

and the latter is often assessed via a χ^2 test of whiteness (see Methods), for each channel $i = 1, \dots, n$. In equation (3), $\varepsilon(t)$ is the same one-step-ahead prediction error in equation (2), and \bar{y}_i is the temporal average of $y_i(t)$ (often equal to zero due to mean centering) and corresponds to a constant predictor which always predicts $\mathbf{y}(t)$ equal to its average $\bar{\mathbf{y}}$. Therefore, it is clear that R_i^2 is always less than or equal to 1 but *can be negative*. A value of $R_i^2 = 1$ indicates a perfect model (for channel/region i), $R_i^2 = 0$ indicates a model as good as the constant predictor, and $R_i^2 < 0$ indicates a model worse than the constant predictor.

Linear models provide maximum prediction accuracy with minimum computational complexity. In this work, we fit and compare several families of linear and nonlinear models, as described below (see Methods for details). We fit each family of models to the data for each participant, thereby finding the optimal model at the global or local level (if the corresponding optimization algorithm is convex or non-convex, respectively). We then compare the resulting best models in each family in terms of their cross-validated fit to held-out data of the same participant (see Methods). The most important ground for comparison is the accuracy of their fit, measured by R_i^2 according to equation (3).

First, consider the results for the fMRI data (**Fig. 2a-b**). While we describe the results obtained using a relatively coarse parcellation here, similar results also hold for finely-parcellated and unparcellated data (cf. **Supplementary Figs. 12-14**). Overall, linear models that directly fit the BOLD time series without (de)convolving with an hemodynamic response function (HRF), either with dense or sparse effective connectivity, and with or without higher-order auto-regressive lags, achieve the highest R^2 . Among nonlinear models, the manifold-based locally-linear model achieves a comparable R^2 . Yet, upon closer inspection of this model, we observe that its window size (which is chosen optimally, see Meth-

ods and **Supplementary Fig. 7**) is very large, effectively making it a globally-linear model. The lack of nonlinearity becomes even clearer when examining the pairwise models. Here, we see that a simple linear model performs as well as the minimum mean squared error (MMSE) model, or even slightly better (**Fig. 2b**, right panel) due to the numerical errors of distribution estimation. We thus infer that the former achieves the highest prediction accuracy achievable by *any* generally nonlinear model, albeit for pairwise prediction.

The second ground for comparison is the whiteness of model residuals, also in held-out data, which indicates that all the dynamics in the data are captured by the model and have not leaked into the residuals (**Fig. 2c-d**). Here, linear models also score higher than nonlinear models, with autoregressive (AR) models clearly outperforming others. However, it is noteworthy that the null hypothesis of whiteness is rejected for the residuals of all methods ($Q/Q_{\text{thr}} > 1$), suggesting the presence of unexplained variance left by all models. Generally, the number of lags and sparsity patterns have little effect on the prediction accuracy of linear AR models for rsfMRI data, a positive but weak effect on the whiteness of the residual, and a negative effect on the computational complexity (**Supplementary Fig. 2**). Similar to the comparison of R^2 values, the only nonlinear model whose whiteness of residuals is comparable to the linear ones is the manifold-based locally-linear model which, as explained above, is effectively linear at the global scale. Also as before, the pairwise linear models achieve a degree of whiteness that is almost identical to the pairwise MMSE estimator, ensuring their optimality among all linear and nonlinear pairwise predictors.

Third and finally, we can compare the models by considering the total time that it takes for their learning and prediction (**Fig. 2e**). When comparing the most efficient linear and nonlinear models, we find that linear models take at least one order of magnitude less time to fit than nonlinear models, as expected. However, linear methods can also be extremely complex to learn; linear models with states at the neural level ('Linear w/ HRF') require the most time to learn due to their high flexibility. Notably, this additional complexity of the 'Linear w/ HRF' or nonlinear methods is not counterbalanced by any benefits in their accuracy or whiteness of residu-

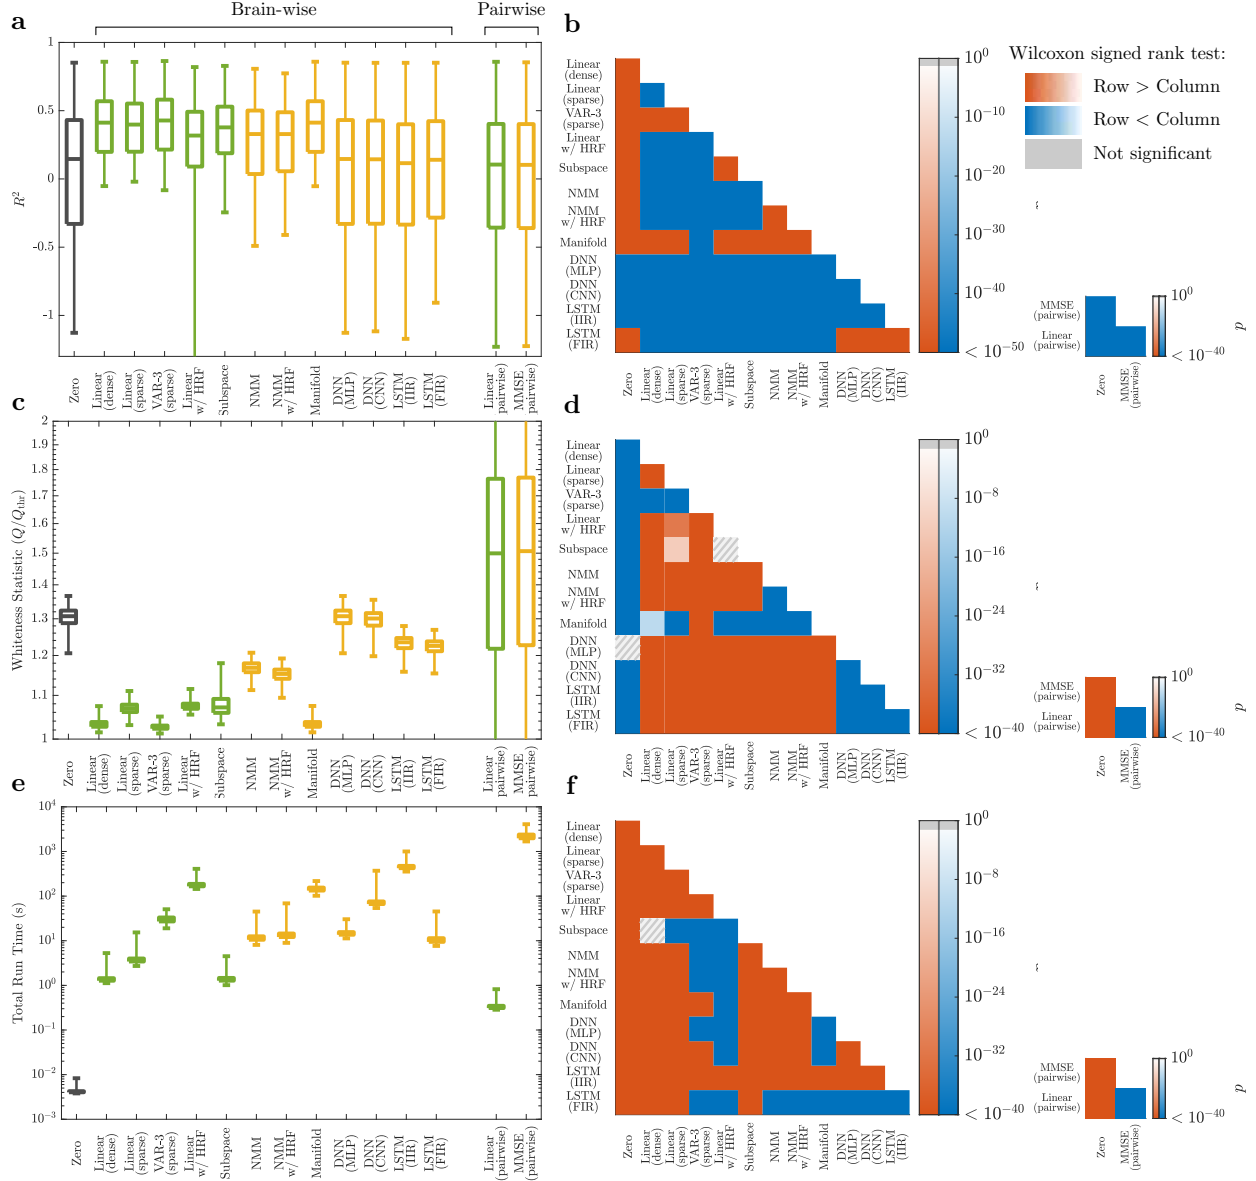


Fig. 2 | Linear vs. nonlinear models of rsfMRI activity. (a) The distribution of cross-validated regional R^2 , combined across all 116 regions and 700 participants, for linear (green) and nonlinear (yellow) models. The gray box corresponds to the zero model used as a baseline (see Methods for an explanation of each model). Negative values of R^2 indicate that models have a worse prediction performance than a constant predictor which always predicts the next value of a signal to be equal to the signal's mean. Note the higher accuracy of the best linear model ('VAR-3 (sparse)') compared to all nonlinear ones. (b) The p -value of one-sided Wilcoxon signed rank test performed between all pairs of brain-wise distributions (left) and all pairs of pairwise distributions (right) of R^2 in panel (a). Warm (cold) colors indicate that the distribution labelled on the row has significantly larger (smaller) samples than the distribution labelled on the column. Gray hatches indicate non-significant differences at an $\alpha = 0.05$ with BH-FDR correction for multiple comparisons. (c, d) Similar to panels (a, b) but for the statistic Q of the multivariate test of whiteness relative to its rejection threshold Q_{thr} (cf. Methods). Smaller Q/Q_{thr} indicates whiter (better) residuals, with $Q/Q_{thr} \leq 1$ required for the null hypothesis of whiteness not to be rejected. (e, f) Similar to panels (a, b) except for the time that it took for the learning and out of sample prediction of each model to run, per participant per cross validation (see Methods). In all box plots, the center line, box limits, and whiskers represent the median, upper and lower quartiles, and the smallest and largest samples, respectively. VAR = vector autoregressive, HRF = hemodynamic response function, NMM = neural mass model, DNN = deep neural network, MLP = multi-layer perceptron, CNN = convolutional neural network, LSTM = long short-term memory, IIR = infinite impulse response, FIR = finite impulse response, MMSE = minimum mean squared error.

als, making the simplest linear models the preferred choice across all measures.

Next, we perform the same comparisons between linear and nonlinear models, but now on the basis of their fit to resting state *iEEG* field potential dynamics (**Fig. 3**). Similar to rsfMRI data, linear AR models provide the best fit to the data, in terms of both the magnitude and whiteness of their cross-validated prediction error. These models also have lower computational complexity than nonlinear ones, with about an order of magnitude (or higher) advantage in computation time.

Alongside these similarities between the rsfMRI and rs*iEEG* data, two major distinctions are notable. First, the R^2 values are generally much higher for *iEEG*, as evidenced by the R^2 distributions of the zero model between the two cases. This difference is due to the fact that the *iEEG* time series has a much higher sampling rate and is therefore smoother. As a result, even predicting each sample equal to its previous sample (i.e., the zero model) has a median accuracy of more than 97% (see **Supplementary Figs. 17-19** for a more detailed assessment of the effects of sampling rate on models' R^2). This fact only highlights the importance of the zero model; without it, the R^2 of all models might have seemed satisfactorily high. In comparison to the zero model, however, it becomes clear that a simple 1-lag linear model, for example, has in fact a very low predictive power.

The second major distinction between the two modalities is the amount of history and temporal dependency within them. fMRI data is almost Markovian, so that $\mathbf{y}(t-1)$ contains almost all the information available for the prediction of $\mathbf{y}(t)$. Little information is also contained in $\mathbf{y}(t-2)$, but almost no information is contained in time points further in the past (**Fig. 2a**). When considering *iEEG* data, in contrast, increasing the number of autoregressive lags up to about 100 still improves the R^2 , although the exact optimal number of lags varies between data segments. In this comparison, it is also important to take into account the vast difference in the sampling frequencies between the modalities, where 2 lags in the fMRI dataset amounts to 1.44 seconds while 100 lags of the *iEEG* data sums to only 0.2 seconds. This greater *richness* of *iEEG* dynamics from a modeling perspective is also responsible, at least in part, for the markedly lower whiteness of residuals of all model families with respect to fMRI (see **Fig. 3c** vs.

Fig. 2c). This greater richness of *iEEG* is also consistent with, though not necessarily a direct consequence of, the fact that *iEEG* data reflects neural signals more directly than fMRI.

The linearizing effects of macroscopic neurodynamics and neuroimaging explain the observed linearity. The above results pose the natural question of why nonlinear models were not able to capture the dynamics in rs-fMRI/rs-*iEEG* data beyond linear ones, even though microscopic neuronal dynamics are fundamentally nonlinear. Here, we focus on four properties of macroscopic neurodynamics and neuroimaging, and show that, in principle, they either fundamentally counteract or apparently mask nonlinearities. Due to its unique position in neural modeling,^{11,12} we will use the sigmoidal nonlinearity to illustrate these effects; we note, however, that the effects are otherwise applicable to other forms of nonlinearity.

The first property that can fundamentally *counteract* microscopic nonlinearities is spatial averaging. Imaging tools that are capable of measuring macroscopic brain dynamics detect a signal that reflects an average over the activity of hundreds, thousands, or even millions of neurons. This spatial averaging can weaken, rather quickly, the nonlinear relationships in the dynamics of individual units (neurons or small-scale neuronal populations) as long as the units are not perfectly correlated, and can completely nullify nonlinearities when correlations decay with distance (**Fig. 4a-c**). Note that this distance can be the physical distance between the units, as assumed here, or in any relevant space such as that of neural codes and stimulus preference. The key factor in the linearizing effect of spatial averaging is the decay of pairwise correlations between neurons so that not all pairs of neurons in a region are highly correlated (a state of blanket global synchrony).

This linearizing effect of spatial averaging is similar to, but different from, stochastic linearization (a.k.a., quasi-linearization).²⁶ While the latter *approximates* the relationship $y = \sigma(x)$ using its expected slope $E[\partial y / \partial x]$, spatial averaging as discussed here can result in a relationship that is truly linear. Also, the same effect can be observed when averaging other forms of nonlinearity than the sigmoid. **Extended Data Fig. 2** shows the effect of spatial averaging on spiking neurons evolving according to the Izhikevich

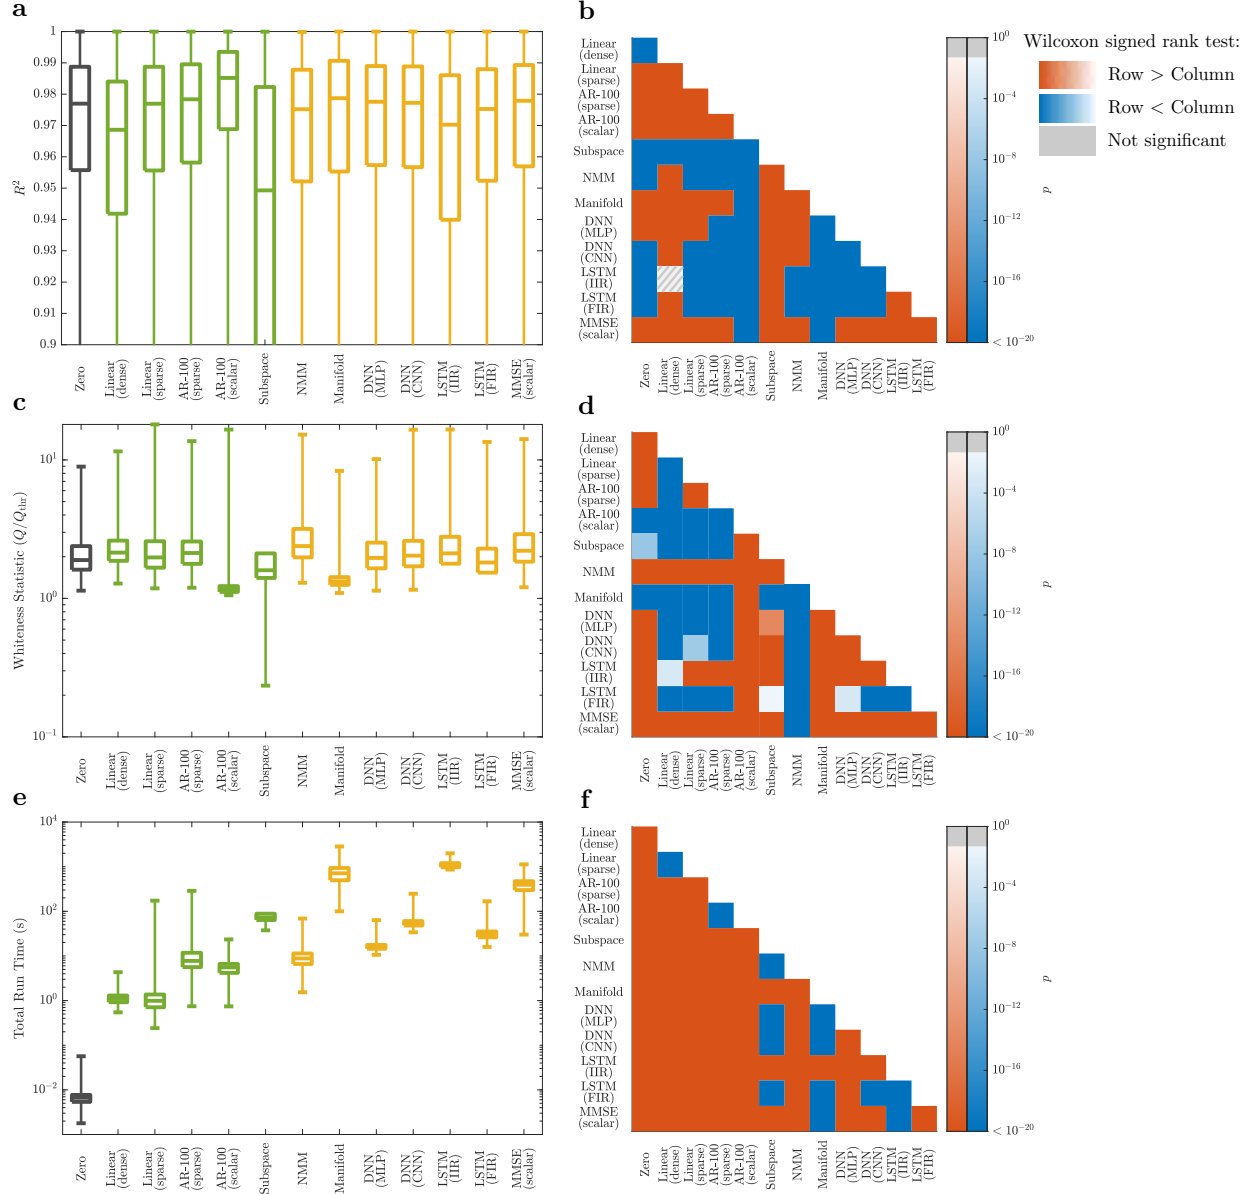


Fig. 3 | Linear versus nonlinear models of rsfMRI activity. Panels and acronyms parallel those in Fig. 2. (a) The distribution of cross-validated regional R^2_i , combined across all electrodes (the number of which varies among participants) and all the recording sessions of the 122 participants (sample size = 776484). Linear and nonlinear methods are depicted by green and yellow boxes, respectively (see Methods for an explanation of each model). Unlike data presented in Fig. 2, pairwise linear or pairwise MMSE models are not included due to the observation that between-electrode connections decrease the cross-validated accuracy of the top model (cf. the 4th and 5th box plots). In contrast, including scalar autoregressive lags is highly beneficial in iEEG, whereas it is not so in rsfMRI. Therefore, the MMSE model here is scalar, conditioning on the past lags of each region itself. The lower whisker of the box plots are trimmed to allow for better illustration of the interquartile ranges. (b) The p -value of the one-sided Wilcoxon signed rank test performed between all pairs of distributions of R^2 in panel (a). Warm (cold) colors indicate that the distribution labelled on the row is significantly larger (smaller) than the distribution labelled on the column. Gray hatches indicate non-significant differences evaluated at $\alpha = 0.05$ with BH-FDR correction for multiple comparisons. (c, d) Similar to panels (a, b) but for statistic Q of the multivariate test of whiteness relative to its rejection threshold Q_{thr} (cf. Methods). Smaller Q/Q_{thr} indicates whiter (better) residuals, with $Q/Q_{thr} \leq 1$ required for the null hypothesis of whiteness not to be rejected. (e, f) Similar to panels (a, b) but for the time that it took for the learning and out-of-sample-prediction of each model. In all box plots, the center line, box limits, and whiskers represent the median, upper and lower quartiles, and the smallest and largest samples, respectively.

model.²³ This model has completely different nonlinearities than the sigmoid (polynomial and discontinuous) and shows a robust nonlinear phenomenon (limit cycle). Although more than a few (but still no more than $100\text{-}10^4$) neurons are required, spatial averaging still dissolves the nonlinear aspects of the dynamics, while mostly sparing the linear ones.

The second property capable of completely counteracting microscale nonlinearities is temporal averaging. Macroscopic neural dynamics are often observed, or even defined, through signals that are low-pass filtered versions of micro- and meso-scale variables. The most notable of these is perhaps the BOLD signal captured by fMRI, which can be seen as an observation of neural activity passed through the low-pass filter of the HRF. Similarly, although to a lesser extent, the local field potentials captured by iEEG most strongly reflect the aggregate pyramidal post-synaptic currents,²⁷ which are themselves low-pass filtered observations of spiking activity through synaptic transmission and neuronal membranes’ resistive-capacitive circuit.²⁸ The effect of low-pass filtering, in essence, is temporal averaging, which impacts nonlinearities in a manner that is similar to that of spatial averaging (**Fig. 4e-f**). The parallel of spatial correlations here is the autocorrelation function or its frequency-domain representation, the power spectral density (PSD). Autocorrelation represents how the correlation between adjacent samples of a signal decay with the temporal distance between those samples. As expected, the smaller the bandwidth of the signal (i.e., the faster their PSD decays with frequency before low-pass filtering), the weaker the linearizing effect of low-pass filtering. As a result, stronger low-pass filtering would also be required to completely nullify nonlinear relationships in signals with narrower bandwidth (**Fig. 4d**). The linearizing effect of temporal averaging also holds for deterministic dynamics, albeit with the resulting linear dynamics (post averaging) also being deterministic (**Extended Data Fig. 3**).

A third property that can counteract or mask nonlinearities is noise. Although both process noise and observation (scanner or electrode) noise may have linearizing effects, here we focus only on the latter. As with any neuroimaging time series, various sources of observation noise can affect the fMRI/iEEG time series^{29,30} and, in turn, “blur” nonlinear relationships, even if they exist between the underlying noise-free

BOLD/LFP time series (**Fig. 4g-h**). In fact, when the power of noise reaches the power of the signal (signal to noise ratio, SNR, ~ 1), it can completely mask a nonlinear relationship in the absence of any spatial or temporal averaging. In reality, however, the linearizing effect of observation noise can combine with spatiotemporal averaging, making the $2 \lesssim \text{SNR} \lesssim 14$ that we have in rsfMRI data (**Supplementary Fig. 4**) potentially more than enough to mask any remaining nonlinearities post-spatiotemporal averaging. Ironically, the use of linear filtering to “clean the data” is more likely to further linearize the dynamics of the time series due to temporal averaging effects discussed above, instead of recovering nonlinearities lost due to noise (Supplementary Note 1). Nonlinear post-processing steps, on the other hand, may leave their own potentially nonlinear signatures in the data, but such signatures should not be confused with true nonlinear relationships in the original BOLD/LFP signal. Further, although we let the noise in **Fig. 4g-h** be independent of the signal, as is typically the case for measurement noise, this linearizing effect would still hold if the noise is linearly dependent on the signal.

The fourth and final property that we will discuss is the number of samples required for detecting nonlinear relationships in large dimensions. Let us assume, despite our discussion so far, that a perfect noise-free nonlinear relationship exists between n -dimensional fMRI or iEEG time series and a noise-free sensor can capture it perfectly. When only $N \sim 1000$ data points are available, we find that the manifold-based predictor – which was our most predictive nonlinear method both for fMRI and iEEG – is still unable to predict the nonlinear relationship better than a linear model in $n \sim 40$ dimensions or higher (**Fig. 4i-j**). This loss in the predictive power of this nonlinear predictor with increasing dimensionality can be easily seen from the fact that the smallest mesh, having two points per dimension, requires an exponentially large $N = 2^n$ data points. Indeed, incorporating structural bias into the learning algorithm can arbitrarily reduce this sample complexity *if* the incorporated bias is consistent with the underlying data³¹ (e.g., if one looks for relationships of the form $y = \sigma(x_1 + \dots + x_n)$ in **Fig. 4i-j**). However, using predictors with structural bias can also be arbitrarily misleading if their form of nonlinearity is not consistent with the given data,³² which is one potential reason for the lower

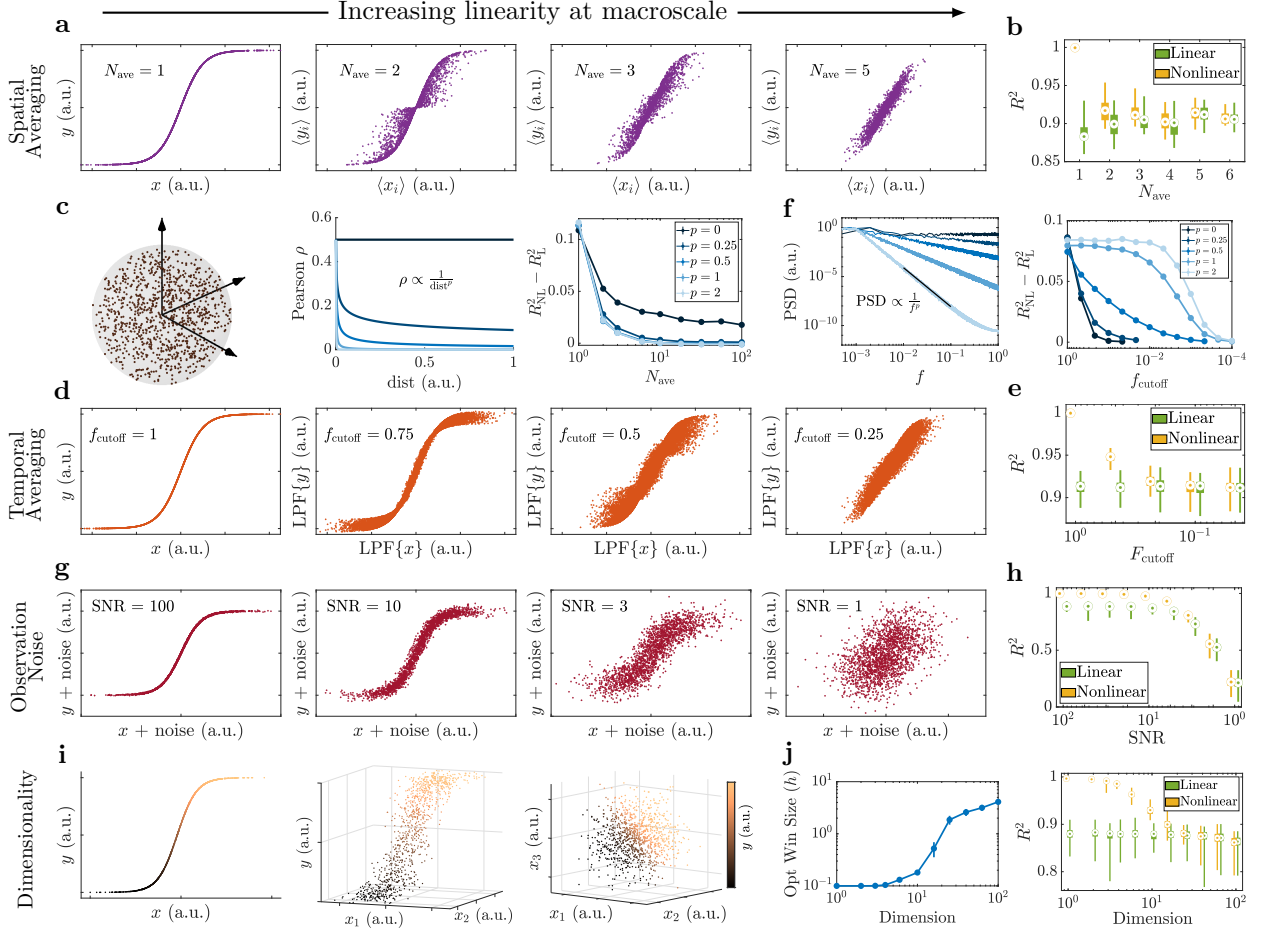


Fig. 4 | The linearizing properties of macroscopic brain dynamics and of neuroimaging measurements. (a) The effect of spatial averaging. For each panel, N_{ave} pairs of signals $x_i(t), t = 1, \dots, 2000$ were randomly and independently generated, $y_i(t) = \tanh(x_i(t))$ was calculated, and their averages $\langle x_i \rangle$ and $\langle y_i \rangle$ were computed. The quantities $\langle x_i \rangle$ and $\langle y_i \rangle$ possess a linear relationship as $N_{\text{ave}} \sim 5$ or higher. (b) The cross-validated R^2 of the optimal nonlinear (MMSE) and linear predictors for the $\langle x_i \rangle$ - $\langle y_i \rangle$ relationships in panel (a). (c) The effect of spatial correlation on spatial averaging. Here, we assign $(x_i(t), y_i(t))$ pairs to spatial locations in a unit sphere (left) and make each $x_i(t)$ and $x_j(t)$ correlated in a manner that depends on their spatial distance (middle). The difference between nonlinear and linear R^2 always decays with N_{ave} and vanishes if the correlation decays, even slowly, with distance (right). (d) The effect of temporal averaging. One pair of $x(t), y(t) = \tanh(x(t))$ is generated, independently over time, and passed through a Gaussian low-pass filter (LPF) with a cutoff frequency f_{cutoff} that is normalized to the Nyquist frequency; thus, $f_{\text{cutoff}} = 1$ means no low-pass filtering. (e) Same as panel (b) but for the LPF $\{x\}$ -LPF $\{y\}$ relationships in panel (d). (f) Similar to panel (c) but for temporal averaging. We varied the PSD decay rate of $x(t)$ (left) and then low-pass filtered $x(t)$ and $y(t) = \tanh(x(t))$ as in panel (d). The difference between the optimal linear and nonlinear R^2 eventually vanishes as f_{cutoff} decreases, but it happens at smaller f_{cutoff} for larger decay rates p . (g) The effect of observation signal to noise ratio (SNR). The quantities $x(t)$ and $y(t) = \tanh(x(t))$ are as in panel (d) and their additive noises are generated independently. (h) Same as panel (e) but for the $(x + \text{noise})$ - $(y + \text{noise})$ relationships shown in panel (g). (i) The effect of dimensionality. The values $x_1(t), \dots, x_n(t)$ are generated as in panel (a) but here $y(t) = \tanh(x_1(t) + \dots + x_n(t))$ generates a one-dimensional nonlinearity in $n + 1$ dimensions. No noise is included; no spatial or temporal averaging is applied. (j) Right: similar to panels (b, e, h) except that a manifold-based (locally-linear) nonlinear predictor is used since the conditional density estimation required for MMSE loses accuracy in high dimensions with fixed number of data points (see Methods). Left: the optimal window size of the manifold-based predictor as a function of dimension n . As n increases, the locally-linear predictor automatically chooses larger windows to be able to make reliable predictions, thereby effectively degrading to a globally linear predictor (see also **Supplementary Fig. 1**). In all box plots, the center point, box limits, and whiskers represent the median, upper and lower quartiles, and the smallest and largest samples, respectively. Error bars in panels (c, f, j) represent one standard error of the mean.

performance of most nonlinear methods in **Figs. 2,3**. This discussion also makes it clear that the inability of our nonlinear system identification methods to outperform linear ones in **Figs. 2,3** over the entire brain is not a proof that no nonlinear method can possibly do so. We can, nevertheless, be certain about this for pairwise or scalar AR models (for fMRI and iEEG, respectively) where the optimal MMSE predictor was computable and performed as well as a linear one.

In conclusion, the process of averaging over space, the process of averaging over time, the existence of observation noise, and the acquisition of limited data are each characteristic of macroscale brain dynamics or neuroimaging measurements, and can transform microscopically nonlinear dynamics into macroscopically linear ones. In reality, their effects are likely all combined, rendering the optimality of linear models in our comparisons not as unexpected as it might originally seem. This linearity has major implications for computational neuroscience, as we discuss next.

Discussion

Summary. In this work, we set out to test the hypothesis that macroscopic neural dynamics are nonlinear, and using linear models for them results in an inevitable loss of accuracy in exchange for simplicity. We thus compared linear and nonlinear models in terms of how well they can predict rsfMRI and rsEEG data in a cross-validated prediction error (PE) system identification framework, where the quality of each model’s fit was assessed by the variance and whiteness of its PE (residual). We found that linear models, and AR models in particular, achieve the lowest PE variance and highest PE whiteness, outperforming neural mass models, deep neural networks, manifold-based models, and the optimal MMSE predictors. In the case of fMRI data, we further verified that the higher predictive power of AR models holds not only in aggregate, but also in a strong majority of individual regions across the brain (**Extended Data Fig. 4a**). Interestingly, the spatial (regional) distribution of the R^2 of the best model also shows significant differences across established cortical functional networks, a remarkably lower predictability of subcortical regions relative to cortical ones, and a close alignment between most methods (all but ‘Linear w/ HRF’) (**Extended Data Fig.**

4b-d). This distinction in predictability highlights significant differences in how spatio-temporally correlated the fMRI time series of different regions are, while the mechanistic physiological and technological reasons behind this distinction remains a warranted avenue for future research.

To further understand the possible causes of the optimality of linear models, we analyzed the effects of common elements of macroscopic neural dynamics: averaging over space and time, observation noise, and limited data samples. We showed that they can each counteract or mask the nonlinearities present at smaller scales. These linearizing effects will add up when combined, suggesting that linear models provide a useful choice for macroscopic neural dynamics at rest; of course, in certain experimental conditions, rigorous system identification methods might still uncover nonlinear dynamics in future studies.

The observed optimality of linear models for the resting state is accompanied by both challenges and opportunities. Having a linear model for neuroscience investigations is computationally ideal, given the extent to which the behavior of linear systems and their response to stimuli are mapped out. Nevertheless, to what extent these linearly-interacting macroscopic signals are informative of, and have a causal influence on, the underlying microscopic activity, remains unclear and represents an invaluable area for future investigation. Our observations also warrant the exploration and development of both linear and nonlinear models of macroscopic neural dynamics beyond those tested here and available in the literature.

Connections to Prior Literature. It is important to distinguish the pragmatic, modeling question that drove our analysis from the rather philosophical question of whether any signs of “nonlinearity” can be found in neuroimaging time series. The latter question has been extensively investigated,^{17–19,33} and often uses determinism or chaos as a proxy for nonlinearity. To answer our distinct modeling question, we used a system identification approach that allows for a direct, side-by-side comparison of linear and nonlinear models. In contrast, the aforementioned studies often resort to indirect, surrogate-based comparisons that rely on strong (and debated) assumptions about the constructed surrogates.³⁴ Also related to, but different from, our work are studies that seek to determine whether the end-to-end input-output mapping

between stimuli and neuroimaging signals (EEG or BOLD) (Ref.^{35,36}) or between functional connectivity and individual phenotypes^{37,38} is (highly) nonlinear. Our focus here, however, is on the nonlinearity of the internal network dynamics of the brain. Finally, past studies have also examined the performance of linear models *per se* in fitting neuroimaging time series (e.g., Ref.³⁹), but without a comparison to nonlinear models.

Our analysis of the linearizing effect of spatial correlation is also related to the large body of work investigating the effect of spatial correlations on the information content and decoding accuracy of neural population codes (see, e.g., Ref.⁴⁰). As expected, the stronger the correlation between neurons, the weaker the linearizing effects of spatial averaging. However, nonlinearities can in principle have two opposing effects on the neural code. On the one hand, nonlinearities can substantially increase the computational complexity and expressivity of a neural network, making correlations beneficial for the neuronal encoding. On the other hand, if the expressivity is too high, the decodability of one neural population by another may decrease, potentially making the linearizing effects of low correlations favorable. Determining which effect dominates, and whether an optimal point exists at the levels of neural correlation observed, remain areas of future research *in vivo*.

Results and Implications. The implications of the linearity of brain dynamics are far-reaching. Linear systems fundamentally have a more limited repertoire of dynamic behaviors than nonlinear ones, excluding the possibility of multi-stability, chaos, limit cycles, or cross-frequency coupling, to name a few.⁴¹ When driven by noise, linear systems act as linear filters that shape the power spectrum of their output (here, fMRI or iEEG time series) through their frequency response, essentially amplifying the frequency content near their resonance frequencies and dampening it elsewhere. Importantly, this effect of shaping the power spectrum of linear systems acts independently over different frequencies; in contrast, nonlinear systems can drive arbitrarily complex cross-frequency interactions.⁴²

The linearity of brain dynamics has even greater implications for network control.^{43,44} The design and analysis of optimal, robust, adaptive, and many other forms of control are much better understood in the

context of linear systems than nonlinear ones. This contrast in tractability only grows for large-scale systems like the brain, thus motivating the recent surge of interest and advancements in using linear control theory in neuroscience.^{45–47} Nonlinear models also present additional challenges beyond network control, including analytical and mechanistic understanding of their functionality, obtaining provable guarantees on their performance, and even hardware requirements for their use in chronic implantable devices. In this context, the present work shows that the favorable tractability and simplicity of linear models do not necessarily come at the often-presumed cost of model inaccuracy, and also provide the necessary tools for identifying the most accurate models for any datasets of interest.

In the analysis of fMRI data, we found that incorporating an HRF component in the model, instead of modeling the dynamics directly at the BOLD level, results in a loss of accuracy in linear models (see ‘Linear (sparse)’ vs. ‘Linear w/ HRF’), and is almost ineffective in nonlinear models (see ‘NMM’ vs. ‘NMM w/ HRF’). It was also in light of this observation that we did not include an HRF component in the majority of our models, such as the DNN or the manifold-based models. This lack of advantage of an explicit HRF component (within the specific context of modeling resting state fMRI dynamics using ODEs) is understandable on a number of grounds. First, in order to include an HRF component in the model, either one should learn the HRF from the data, such as in our ‘linear w/ HRF’ model, which will create marked model flexibility and therefore increase the likelihood of over-fitting; or, one should use a typical HRF, such as in our ‘NMM w/ HRF’ model, which is a source for additional error. Second, by including the HRF in the model, we ultimately seek to recover neural information that is lost through the HRF. This task is difficult, if not impossible, without a high signal-to-noise ratio as well as more accurate HRF models than those currently available. Finally, a linear autoregressive model can automatically capture a linear approximation of the HRF dynamics,⁴⁸ precisely as present in the observed time series. Ultimately, our results encourage a side-by-side comparison of models with and without the inclusion of an HRF component in order to assess the costs and benefits of such inclusion for any datasets of interest.

A very similar argument also applies to including

colored noise in the model, both for fMRI and iEEG. Even though the so-called “noise floor” in neural signals, over which neural oscillations are detected,⁴⁹ has a clear power-law ($1/f$) PSD, its decay can be well modeled by white noise passing through a linear filter. This property of the $1/f$ noise is in fact why the AR models, which assume a white noise signal $\mathbf{e}(t)$, have prediction errors that are maximally white. This latter fact can be directly seen from **Figs. 2c,3c**, when noting that any model’s prediction errors are, by construction, the model’s estimate of the noise $\mathbf{e}(t)$ (Ref.²⁵).

In addition to considering linear and nonlinear models, we wish to underscore the importance of the zero model. It is not uncommon in the modeling literature to assess the quality of a fitted model *per se*, without any grounds for comparison. For instance, our ‘DNN (MLP)’ model for fMRI had a median R^2 of about 14% and for some participants it had a median R^2 (among all regions) of over 50%. Even more notably, the ‘DNN (MLP)’ model for iEEG had a median R^2 of over 97%. Without any comparisons, these numbers may suggest that the ‘DNN (MLP)’ models are quite accurate; yet, as seen in **Figs. 2b,3b**, the predictive accuracy of these models are in fact lower than the zero model in fMRI, and indistinguishable from it in iEEG. The act of comparing to a baseline model is therefore an essential step in the assessment of any model’s goodness-of-fit.

We restricted our analyses here in the main text to certain spatiotemporal resolutions for both fMRI (a coarse parcellation) and iEEG (a high sampling rate), naturally raising the question of how robust our findings are to our choices of resolutions. As shown in **Supplementary Figs. 12-19**, our main finding (higher predictive power of linear autoregressive models over all other model families) holds across all resolutions tested. We do, however, observe certain differences between resolutions. In iEEG data, we observe that using lower sampling rates (and therefore longer time intervals) increases the benefit of modeling network interactions, even whilst lowering the R^2 values across all models. In fMRI data, we interestingly see that as we move towards more fine-grained parcellations and ultimately unparcellated data, (i) the simpler ‘Linear (sparse)’ model with less parameters gains advantage over the more populated ‘VAR-3 (sparse)’, and (ii) the overall R^2 values of all models is reduced, potentially due to the improvement

in signal-to-noise ratio resulting from averaging in coarser parcellations.

Methodological Considerations. Despite the solid theoretical foundations of the PE method for system identification, our results may still beg a practical question: would the same system identification and side-by-side comparison procedure be able to identify nonlinear dynamics, should they actually exist in the time series? A direct answer to this question can be given, e.g., by applying the same procedure to simulated time series generated from a nonlinear model whose ground truth functional form we know. The result of such an analysis is provided directly in **Extended Data Figs. 2,3**, but also indirectly in **Fig. 4**. Note that in the latter, we compared the cross-validated predictive power of linear and (MMSE or manifold-based) nonlinear models in identifying the sigmoidal relationship $y = \sigma(x)$ and its variants, after averaging or noise addition. This relationship can be equally viewed as a nonlinear dynamical system $\dot{x} = \sigma(x)$, the nonlinearity of which was only identified until counteracted or masked by the four macroscopic effects we discussed therein.

In **Supplementary Fig. 9**, we take this validation one step further and compare the ‘DNN (MLP)’ model with linear autoregressive and subspace models while also tuning their respective hyper-parameters using simulated data from the Izhikevic model and the same stochastic gradient descent (SGD) algorithm used for hyper-parameter tuning in the main results. The resulting higher predictive performance of the ‘DNN (MLP)’ model is another testament to the validity of the end-to-end process we used for model training and comparison. Note, also, that the amount of improvement that nonlinear models can offer over linear ones can vary widely based on the true underlying dynamics and the ability of either model to capture them. In **Supplementary Fig. 10**, e.g., we show a similar side-by-side comparison between the ‘DNN (MLP)’ and ‘Linear (dense)’ models on one of the simplest nonlinear systems exhibiting a nonlinear behavior (chaos in this case). Despite the extreme simplicity of the true underlying model, the neural network can achieve near-perfect accuracy with only one hidden layer and 10 hidden units while the linear model can achieve about 50% R^2 . Finally, we also trained both the parameters and hyper-parameters of the ‘DNN (MLP)’ model on fMRI data while re-

placing its ReLU activation functions with linear ones using the same codes and routines used for the actual ‘DNN (MLP)’ model. As seen from **Supplementary Fig. 11**, the resulting model achieves indistinguishable R^2 distributions with the (theoretically equivalent) ‘Linear (dense)’ model, except for implementation differences between our hard-coded linear regression solver and that provided by the MATLAB’s *trainNetwork* function.

As for the comparison of linear and nonlinear models, one might expect nonlinear models to perform at least as well as linear ones, but not worse, given that the space of nonlinear models includes all linear models as a special case. In our comparison, however, we saw that most of our nonlinear methods actually have a worse prediction performance than linear ones. This behavior can be understood in light of at least two facts. First, many nonlinear models, such as neural mass models, do not include linear models as a special case and have structural biases that can be a source of error if not consistent with the data.³² Second, even nonlinear models that do not have structural biases and/or contain linear models as a special case, such as DNN, MMSE or manifold-based models, still have a marked flexibility relative to a linear model. This immense flexibility makes the training of these models and finding their global optimum challenging. As a result, training algorithms are quite likely to return suboptimal models which, in this case, show worse accuracy and generalization than their linear special cases.

A further noteworthy aspect of our study specifically, and the prediction error framework more generally, is the focus on fitting the time series rather than its derivative statistics, such as the functional connectivity (FC) (Ref.^{50–52}) or power spectral density.⁵³ While the choice of one approach over the other ultimately depends upon the anticipated use of the learned model, it is important to note that the mapping from dynamical systems to FC (or any other such statistic) is not a one-to-one mapping.⁵⁴ In fact, linear systems of the form $\mathbf{y}(t) - \mathbf{y}(t-1) = \mathbf{W}\mathbf{y}(t-1) + \mathbf{e}(t)$ with completely different \mathbf{W} matrices can give rise to almost identical FC matrices (see **Extended Data Fig. 1**). Therefore, when considering the accuracy of a general purpose model of the brain, the time series contains the maximum amount of information and thus provides the best target for model fitting.

One modeling approach that we did not employ in this study is dynamic causal modeling (DCM, Ref.⁵⁵). The reason is that neither of the current variants of DCM are feasible, due to their computational complexity at the scale of our analysis: whole brain fMRI with $n = 116$ parcellations or large-scale iEEG with up to 175 and a median of 98 electrodes. The most efficient variant, spectral DCM for fMRI, for instance, is applicable to ~ 30 -40 nodes, whereas stochastic DCM (the most relevant to our study) is only applicable to much smaller systems. However, in light of our results thus far, the great computational complexity of the DCM approach, and thus its potential for overfitting, we would not expect its cross-validated R^2 to reach that of a linear model, although this comparison remains unknown at the present.

In this work, we demonstrated four properties of macroscopic neurodynamics that can counteract or mask microscopic nonlinearity. In doing so, we purposefully kept the discussion at a conceptual level and generally abstained from tying it to specific micro- or mesoscopic neural models, as doing so would require building on assumptions that our study explicitly seeks to avoid. For instance, it is currently unclear whether, and to what extent, the dynamics of the *mesoscopic* local field potentials or population firing rates that seem to be the main neural drivers of fMRI or iEEG are nonlinear and, if so, what the precise form of their nonlinearity is at each brain region. A warranted avenue for future research would be the re-analysis of the effects of spatial and temporal averaging, observation noise, and limited data samples on precise, data-driven models of mesoscopic brain dynamics, should they possess nonlinear interactions.

Limitations. Finally, we would like to highlight some of the limitations of the present study. First, it is important to note that the space of all nonlinear models in tens to hundreds of dimensions is intractably large, and the fact that our tested nonlinear models did not outperform linear ones is not a proof that no nonlinear model may ever do so.

In particular, some nonlinear models (e.g., ‘DNN (MLP)’) have the capacity to implement the ‘Zero’ or even linear AR models as a special case (e.g., by setting the output weights of the ‘DNN (MLP)’ model to zero to replicate the ‘Zero’ model or to adjust its weights and biases so that all ReLU activation functions operate in their linear range to replicate

the ‘Linear (dense)’ model). Therefore, despite our best efforts in training these models and validating our training algorithms in alternative simulated settings (cf. **Supplementary Figs. 7-9**), the fact that their predictive power does not reach the performance of the zero/linear models in some datasets suggests suboptimal parameter and/or hyperparameter fitting due, e.g., to local minima or a lack of enough training data. Therefore, we should highlight that our achieved predictive power of those nonlinear methods do not necessarily reflect their best performance and it is possible that future work with more training data and/or alternative optimization approaches may get better results from these nonlinear models that match or outperform linear models in iEEG/fMRI.

Our work thus seeks to provide rigorous evidence and methodology towards resolving the linear vs. nonlinear modeling dilemma in computational neuroscience, rather than a final resolution thereof. Hybrid models consisting of various combinations of different nonlinear structures (LSTMs and CNNs, for example), may have substantial potential for obtaining more accurate nonlinear models and warrant future research. We can be confident, nevertheless, about the optimality of linear models at the pairwise level for fMRI or scalar AR level for iEEG given the equal or higher prediction power of linear regression relative to the optimal MMSE predictor. Moreover, our modeling framework is currently only applicable to resting state dynamics with no inputs, and has been tested on the two modalities of fMRI and iEEG. Inclusion of input signals for system identification of task fMRI/iEEG data requires accurate data-driven ‘input models’ of how experimental stimuli, as well as participants’ voluntary responses, influence the BOLD or LFP signals in each brain region, and is a highly warranted avenue for future research.^{56,57} Under intensive task conditions, moreover, it is more likely, or perhaps certain, to observe nonlinearities at least in the form of saturation effects in the BOLD/LFP signal. However, the precise form and extent of this nonlinearity needs to be determined using rigorous system identification routines.

In conclusion, our work sought to ask the often unasked question of whether the brain is macroscopically linear. Our findings show that simple linear models explain the rsfMRI and rsiEEG data as well as, or even better than, an array of nonlinear ones, thus challenging the commonly held, yet untested

assumption of higher accuracy of nonlinear models. However, the costs and benefits of nonlinear models are ultimately case-specific. Therefore, instead of offering a universal recommendation on the preferable choices for the modeling of neural dynamics, we rather provide the groundwork for rigorous investigation and informed decision-making in the context of rsfMRI/rsiEEG. When feasible, following a similar system identification routine is always recommended for computational modeling of any datasets of interest, in order to ensure the optimal fit of the models used for subsequent analysis or design.

Methods

Data and pre-processing. For the fMRI analysis, we used ICA-FIX resting state data from the S1200 Human Connectome Project release.^{58,59} The HCP experiments were carried out by the WU-Minn consortium and its adherence to ethical standards was approved by the by the Internal Review Board of the respective institutions. Explicit informed consent was acquired from all participants involved in the study. rsfMRI images were collected with the following parameters: TR = 720 ms, TE = 33.1 ms, flip angle = 52 deg, FOV = 208x108 mm, matrix = 104x90, slice thickness = 2.0 mm, number of slices = 72 (2.0 mm isotropic), multi factor band = 8, and echo spacing = 0.58 ms. Brains were normalized to fsLR32k via the MSM-AII registration and global signal was removed. No bandpass filtering was performed (see Supplementary Note 1). Finally, we removed participants from further analysis if any of their four resting state scans had excessively large head motion, defined by having frames with greater than 0.2 mm frame-wise displacement or a derivative root mean square (DVARS) above 75. Also, participants listed in⁶⁰ under “3T Functional Preprocessing Error of all 3T RL fMRI runs in 25 Subjects” or “Subjects without Field Maps for Structural scans” were removed, leaving a total of 700 participants that were used for all the analyses. We parcellated the brain into 100 cortical regions (Schaefer 100x7 atlas⁶¹) and 16 subcortical ones (Melbourne Scale I atlas⁶²).

For iEEG preprocessing, raw data from the RAM data set we have published on previously⁶³⁻⁶⁵ was segmented into task free epochs from either before or after task completion that were at least 5 min-

utes in length. This process resulted in a total of 283 recordings from 122 participants. Data were then downsampled to the lowest sampling rate used across recording sites (500 Hz). Electric line noise and its harmonics at 60, 120, and 180 Hz were filtered out using a zero phase distortion 4-th order stop-band Butterworth filter with a 1 Hz width. This procedure was implemented using the *butter()* and *filtfilt()* functions in MATLAB. We then rejected noisy channels that were either (i) marked as noisy in the RAM dataset notes, (ii) had a line length greater than three times the mean, (iii) had z -scored kurtosis greater than 1.5, or (iv) had a z -scored power-spectral density dissimilarity measure greater than 1.5. The dissimilarity measure used was the average of one minus the Spearman’s rank correlation with all channels. Data were then demeaned and detrended. Channels were grouped according to whether they were grid or depth electrodes, and then common average referenced within each group. Following the common average referencing step, plots of raw data and power spectral densities were visually inspected by an expert researcher with 6 years of experience working with electrocorticography data to ensure that data were relatively clean.

Finally, while the aforementioned channel removal criteria are consistent with the standard practice in iEEG pre-processing (where we remove channels with exceptional line length, kurtosis, and power spectral densities to target high-frequency noise and ictal activity,^{66,67} electrode drift and ictal spikes,^{68–70} or line noise and flat power spectral densities,^{65,71} respectively) and they are essential from a data quality perspective, see also **Supplementary Fig. 22**. Here, we reproduced our main finding of the higher predictive power of linear models on data without channel removal as a validation that channel removal had not confounded our finding by potentially providing an edge for the linear models.

Computing and run time calculations. All the computations whose run time was measured and reported in **Figs. 2e,3e** were performed on the CUBIC cluster at the University of Pennsylvania, using 1 CPU core and 16 or 64 GB of memory per fMRI or iEEG computing jobs, respectively. For each method, fMRI participant/iEEG segment, and cross-validation fold, one training time, one test time, and one total time was computed, where the latter is sim-

ply the sum of the former two. Note that for the ‘Zero’, ‘Manifold’, and the ‘MMSE’ models no training time can be defined. This should be clear for the ‘Zero’ model, but it is also the case for the ‘Manifold’ and the ‘MMSE’ models due to their “model on demand” nature, i.e., that all training data is directly used in computing the prediction of each test point, instead of the usual process of learning one model from the training data and then using that model for test predictions. Therefore, the training time is set to zero for these models. For the remaining models, the training time equals the CPU time taken for all the computations from the point where data has been broken into train and test and until the point where the model is learned. The test time, well-defined and non-zero for all the methods, starts immediately after the end of the training period and runs until the point where the output time series (one-step ahead predictions) are computed. Note that computations of the R^2 and whiteness statistics are optional post-hoc analyses and done in the same way for all methods, hence not included in any of these run time calculations. A breakdown of the run time into separate training and test times is also shown in **Supplementary Figs. 5,6**.

Linear and nonlinear families of models. A complete list of the models used in this study are provided in **Supplementary Table 1**. For all models, the continuous-time dynamics in equation (1) is first discretized. With a slight abuse of notation, we also represent the discretized dynamics as

$$\mathbf{x}(t) - \mathbf{x}(t-1) = f(\mathbf{x}(t-1)) + \mathbf{e}_1(t), \quad (4a)$$

$$\mathbf{x}(0) = \mathbf{x}_0, \quad t = 1, \dots, N$$

$$\mathbf{y}(t) = h(\mathbf{x}(t)) + \mathbf{e}_2(t), \quad t = 0, \dots, N \quad (4b)$$

where the time index t is now an integer, for simplicity of notation, but the discretization step size is always equal to 720ms seconds for the HCP data (equivalent to 1 TR) and 2ms for the RAM data. This choice means that, e.g., the map f in equation (4a) equals 1 time step multiplied by the map f in equation (1a), and $T = 864s$ in equation (1a) corresponds to $N = 1200$ in equation (4a) for the fMRI data. Recall that in this general form, the noise signals $\mathbf{e}_1(t)$ and $\mathbf{e}_2(t)$ can have arbitrary statistics, including white or colored PSD. We then learn the dynamics in equation (4) using the following families of models.

The hyper-parameters used for each model are listed in **Supplementary Table 1** (see ‘Hyper-parameter selection’ below and **Supplementary Figs. 7,8** for details).

Linear models with states at the BOLD/LFP level (‘*Linear (dense)*’, ‘*Linear (sparse)*’): This model is our simplest. In it, we let $\mathbf{y}(t) = \mathbf{x}(t)$, modeling the dynamics directly at the BOLD/LFP level. This also allows for combining the noise signals $\mathbf{e}_1(t)$ and $\mathbf{e}_2(t)$ into a single noise signal $\mathbf{e}(t)$, which is then taken to be white. These simplify equation (4) to

$$\mathbf{y}(t) - \mathbf{y}(t-1) = f(\mathbf{y}(t-1)) + \mathbf{e}(t). \quad (5)$$

If we further let $f(\mathbf{y}(t)) = \mathbf{W}\mathbf{y}(t)$ be linear, then we get

$$\mathbf{y}(t) - \mathbf{y}(t-1) = \mathbf{W}\mathbf{y}(t-1) + \mathbf{e}(t) \quad (6)$$

where \mathbf{W} is an n -by- n matrix of effective connectivity between brain regions. We fit and compare this model both when \mathbf{W} is dense and when it is sparse. The latter is motivated by the facts that (i) from a mechanistic perspective, an important property of brain networks and other large-scale complex networks is their sparsity; while (ii) from a machine learning perspective, regularization and reducing the number of free parameters in a model can prevent over-fitting and improve generalization. To promote sparsity, we use standard 1-norm (LASSO) regularization with a λ hyper-parameter that is tuned separately for fMRI and iEEG.

Linear autoregressive models (‘*AR-2 (sparse)*’, ‘*VAR-2 (sparse)*’, ‘*AR-3 (sparse)*’, ‘*VAR-3 (sparse)*’, ‘*AR-100 (sparse)*’, ‘*AR-100 (scalar)*’): motivated by the long history of AR models in neuroscience,^{24,72,73} here we extend equation (6) to

$$\begin{aligned} \mathbf{y}(t) - \mathbf{y}(t-1) = & \mathbf{W}\mathbf{y}(t-1) + \mathbf{D}_2\mathbf{y}(t-2) \\ & + \mathbf{D}_3\mathbf{y}(t-3) + \dots \\ & + \mathbf{D}_d\mathbf{y}(t-d) + \mathbf{e}(t) \end{aligned} \quad (7)$$

for an ‘AR- d ’ model. The number of lags d was tuned separately for fMRI and iEEG, and the matrix \mathbf{W} is either made sparse using LASSO or enforced to be diagonal. Note that the latter results in n scalar AR models at each node which are completely decoupled from each other. We restricted the matrices $\mathbf{D}_2, \mathbf{D}_3, \dots$ to be diagonal in ‘AR’ models but not so in full vector auto-regressive (‘VAR’) models. In both

cases, we use LASSO regularization to promote sparsity in the regressors, signified by the ‘(sparse)’ suffix in method identifiers, with the regularization hyper-parameter λ chosen optimally and separately for each model (cf. ‘Hyper-parameter selection’ below). In general, we found that λ is a moderately sensitive parameter, more so for the whiteness of residuals than R^2 (cf. **Supplementary Fig. 3** for an example).

Linear models with states at the neural level (‘*Linear w/ HRF*’, only applicable to fMRI data): A standard step in the computational modeling of fMRI dynamics is to incorporate a model of the hemodynamic response function (HRF), and to separate the underlying neuronal variables from the observed BOLD signals. In this family of models, we thus separate the states \mathbf{x} from the outputs \mathbf{y} , while keeping a one-to-one relationship between the two ($m = n$). We then let the latter be a filtered version of the former through the HRF. For generality and given the natural and important variability of HRF across the brain,^{74,75} we allow the HRF to vary regionally and learn it from the data for all regions in addition to the effective connectivity matrix \mathbf{W} . Furthermore, for the sake of generality, we allow both $\mathbf{e}_1(t)$ and $\mathbf{e}_2(t)$ to be colored, with power spectral densities that can also be different between regions and are learned from data. Note that this choice includes, as a special case, white $\mathbf{e}_1(t)$ and $\mathbf{e}_2(t)$. The result is a highly flexible linear model given by

$$\mathbf{x}(t) - \mathbf{x}(t-1) = \mathbf{W}\mathbf{x}(t-1) + \mathcal{G}_1(q)\hat{\mathbf{e}}_1(t) \quad (8a)$$

$$\mathbf{y}(t) = \mathcal{H}(q)\mathbf{x}(t) + \mathcal{G}_2(q)\hat{\mathbf{e}}_2(t)$$

$$\mathcal{H}(q) = \sum_{p=1}^{n_h} \text{diag}(\mathbf{H}_{:,p})q^{-p} \quad (8b)$$

$$\mathcal{F}_1(q) = \mathbf{I} - \mathcal{G}_1^{-1}(q) = \sum_{p=1}^{n_\phi} \text{diag}(\mathbf{\Phi}_{:,p})q^{-p} \quad (8c)$$

$$\mathcal{F}_2(q) = \mathbf{I} - \mathcal{G}_2^{-1}(q) = \sum_{p=1}^{n_\psi} \text{diag}(\mathbf{\Psi}_{:,p})q^{-p} \quad (8d)$$

Since LASSO regression produced the best results in our BOLD-level linear models, we use LASSO (with the regularization weight λ) to promote sparsity in \mathbf{W} here. $\mathcal{H}(q)$ is a diagonal matrix whose (i, i) entry is a linear finite-impulse response (FIR) approximation of the HRF in region i , parameterized as in equation (8b) (q^{-1} is the standard delay operator, such that $q^{-1}x(t) = x(t-1)$, see²⁵). Similarly, \mathcal{G}_1 and \mathcal{G}_2 are diagonal filters, parameterized by the inverse FIR forms in equation (8c)-(8d). The matrices $\mathbf{H}_{n \times n_h}$, $\mathbf{\Phi}_{n \times n_\phi}$, and $\mathbf{\Psi}_{n \times n_\psi}$ include learnable FIR parameters of $\mathcal{H}(q)$, $\mathcal{G}_1^{-1}(q)$, and $\mathcal{G}_2^{-1}(q)$, respectively. Since

the state vector $\mathbf{x}(t)$ is not measured, we learn this model by iterating between state estimation and parameter estimation in an expectation-maximization (EM)-like manner. Note that the presence of filters increases the effective state dimension of the system to $n \cdot \max\{n_\phi + 1, n_h + n_\psi\}$, significantly increasing the computational complexity of the state estimation step. The final model is taken from the EM iteration with the highest (training) R^2 .

Linear models with abstract data-driven states ('Subspace'): The previous model, despite and because of its extreme generality and flexibility, has a very large state dimension and is extremely difficult to fit. If we forgo the physiological interpretability of the states, then significantly simpler and lower-dimensional models of the form

$$\begin{aligned}\mathbf{x}(t) - \mathbf{x}(t-1) &= \mathbf{W}\mathbf{x}(t-1) + \mathbf{e}_1(t) \\ \mathbf{y}(t) &= \mathbf{C}\mathbf{x}(t) + \mathbf{e}_2(t) \\ \text{Cov} \left(\begin{bmatrix} \mathbf{e}_1(t) \\ \mathbf{e}_2(t) \end{bmatrix} \right) &= \begin{bmatrix} \mathbf{Q} & \mathbf{M} \\ \mathbf{M}^T & \mathbf{R} \end{bmatrix}\end{aligned}$$

can be learned via subspace identification method.²⁵ Unlike the model above, states represent abstract low-dimensional regularities within the data, with a dimension n_x that is chosen optimally for each data type. The hyper-parameters r and s represent the amount of output and input-output history used, respectively, in the output Hankel matrix construction and subspace projection steps of the algorithm, respectively (see Ref.²⁵ Ch. 10 for details). The noise sequences $\mathbf{e}_1(t)$ and $\mathbf{e}_2(t)$ are assumed to be white but can be correlated, and the covariance matrices \mathbf{Q} , \mathbf{M} , and \mathbf{R} are also learned from data. Note that this whiteness assumption on noise sequences is without loss of generality in this case due to the subspace method's ability to learn any non-white dynamics (i.e., color) of noise as part of the abstract state dynamics.

Nonlinear neural mass models ('NMM', 'NMM w/ HRF'): Learning of the models above, except for the 'Linear w/ HRF', involves a convex optimization that can be efficiently solved to find its unique global optimum. In contrast, the learning of nonlinear models is less straightforward. Recently, Ref.⁷⁶ developed an algorithm called MINDy that uses state-of-the-art optimization techniques for learning a neural mass

model of the form

$$\begin{aligned}\mathbf{x}(t) - \mathbf{x}(t-1) &= (\mathbf{W}\psi_\alpha(\mathbf{x}(t-1)) - \mathbf{D}\mathbf{x}(t-1))\Delta_T \\ &\quad + \mathbf{e}_1(t) \\ \mathbf{y}(t) &= \mathcal{H}(q)\mathbf{x}(t) + \mathbf{e}_2(t)\end{aligned}$$

using rsfMRI data. In this model, $\mathbf{x}(t)$ has the same dimension as $\mathbf{y}(t)$ (one neural mass per brain region), Δ_T is the sampling time, \mathbf{W} is a sparse connectivity matrix, \mathbf{D} is a diagonal self decay matrix, $\psi_\alpha(\cdot)$ is an element-wise sigmoidal nonlinearity whose steepness is determined by each element of the vector α (which is also the same size as \mathbf{x}), and $\mathcal{H}(q)$ is a scalar linear HRF that is the same and fixed *a priori* for all regions. The associated toolbox that we use allows the user to either deconvolve $\mathbf{y}(t)$ using a canonical HRF to obtain the state $\mathbf{x}(t)$ ('NMM w/ HRF'), or set $\mathcal{H}(q) = 1$ and directly fit the model to $\mathbf{y}(t)$ ('NMM'). We use both methods for fMRI data but only the latter for iEEG. Since the MINDy algorithm was originally tuned for fMRI, we re-tune its regularization hyper-parameters $\lambda_1, \dots, \lambda_4$ for use with iEEG data (see 'Hyper-parameter selection' below).

Nonlinear models via multi-layer perceptron deep neural networks ('DNN (MLP)'): Here we use a model of the form in equation (5) for fMRI and train a rectified linear unit (ReLU) MLP DNN to approximate the function $f(\cdot)$. The structure of the DNN consists of an input layer, D ReLU layers, each preceded with fully connected and batch normalization layers and succeeded with a 50% dropout layer, a final fully connected layer, and the output layer. Given the importance of AR lags in the modeling of iEEG, for this modality we generalize equation (5) as

$$\mathbf{y}(t) - \mathbf{y}(t-1) = f(\mathbf{y}(t-1), \dots, \mathbf{y}(t-d)) + \mathbf{e}(t) \quad (9)$$

and similarly approximate $f(\cdot)$ using an MLP DNN. We use MATLAB's Deep Learning Toolbox for the training and evaluation of the DNN and tune the depth D and width W of the DNN separately for fMRI and iEEG (see 'Hyper-parameter selection' below).

Nonlinear models via convolutional deep neural networks ('DNN (CNN)'): Given the recent success of CNNs in complex learning problems, we also included a model similar to 'DNN (MLP)' but with a CNN to approximate the function $f(\cdot)$. The network consists of an input layer, D one-dimensional convolutional

layers (convolving over time using n_{filt} filters of size l_{filt}) each succeeded by a batch normalization layer, a ReLU layer, and an average pooling layer with a pool size of n_{pool} , a final dropout layer with probability p_{drop} , a fully connected layer, and the output layer. Spatial convolution was not included in the model, as is the standard in modeling dynamical systems with CNNs, due to the arbitrary nature of channel numbering. Temporal convolution is nevertheless the basis of this model and we thus consider $d > 1$ autoregressive lags for both fMRI and iEEG.

Nonlinear models via long short-term memory neural networks ('LSTM (IIR)', 'LSTM (FIR)'): The above DNN models are inherently static (i.e., feedforward), whereas various recurrent neural network architectures have also been proposed for directly modeling dynamical systems. One of the most successful such architectures are LSTMs which we implemented here in two forms: infinite impulse response (IIR) and finite impulse response (FIR). These two forms correspond, respectively, to the two common sequence-to-sequence and sequence-to-one forms of modeling time series using LSTMs. In both cases, the network consists of an input layer, a layer of W LSTM units, a fully connected layer, and an output layer. The difference is that in the IIR model, the network is initialized once at time 0 and run forward, continuously receiving $\mathbf{y}(t-1)$ as input and generating $\mathbf{y}(t) - \mathbf{y}(t-1)$ as output. Each output, therefore, depends on the entire history of the inputs. In the FIR model, on the other hand, the model is initialized and run forward once for each time point t , receiving only $\mathbf{y}(t-d), \dots, \mathbf{y}(t-1)$ as input and predicting $\mathbf{y}(t) - \mathbf{y}(t-1)$ as output.

Nonlinear manifold-based models ('Manifold'): Consider equation (5) or equation (9) and assume, for simplicity, that f is differentiable. Each of these systems of equations consists of n scalar equations, each of which defines a manifold (surface) in $n+1$ -dimensional (for equation (5)) or $nd+1$ (for equation (9)) space. Various methods have been developed in the machine learning and system identification literature^{77,78} and used in computational neuroscience^{20,79} to capitalize on the fact that in the small vicinity of a point, the manifold can be approximated by a linear hyperplane tangential to it at that point. Here, we use the simple method of local polynomial modeling of order 1 (Ref.⁷⁷). To explain this method, first consider the simpler model in equation (5). For

each test time t^ℓ , we approximate the function $f(\cdot)$ as a linear function in the vicinity of $\mathbf{y}(t^\ell - 1)$, i.e.,

$$f(\mathbf{z}) \simeq \mathbf{c}^\ell + \mathbf{W}^\ell [\mathbf{z} - \mathbf{y}(t^\ell - 1)] \quad (10)$$

The constant vector \mathbf{c}^ℓ and matrix \mathbf{W}^ℓ are learned from training data (separately for each test point), as follows. Each training point $\mathbf{y}(t^m - 1)$ is weighted according to its distance to $\mathbf{y}(t^\ell - 1)$, i.e.,

$$k^{\ell,m} = \exp \left(- \frac{\|\mathbf{y}(t^m - 1) - \mathbf{y}(t^\ell - 1)\|^2}{2h^2} \right) \quad (11)$$

where the hyper-parameter h controls how local or global the model is. These weights are then used in a weighted least squares estimation,

$$\begin{aligned} [\mathbf{c}^\ell \quad \mathbf{W}^\ell] &\simeq \left(\sum_{t^m} [\mathbf{y}(t^m) - \mathbf{y}(t^m - 1)] k^{\ell,m} \boldsymbol{\varphi}^{\ell,mT} \right) \\ &\cdot \left(\sum_{t^m} \boldsymbol{\varphi}^{\ell,m} k^{\ell,m} \boldsymbol{\varphi}^{\ell,mT} \right)^\dagger \end{aligned}$$

where † denotes pseudo-inverse and

$$\boldsymbol{\varphi}^{\ell,m} = \begin{bmatrix} 1 \\ \mathbf{y}(t^m - 1) - \mathbf{y}(t^\ell - 1) \end{bmatrix}. \quad (12)$$

Note that only the computed \mathbf{c}^ℓ is ultimately used, since substituting $\mathbf{z} = \mathbf{y}(t^\ell - 1)$ in equation (10) gives

$$f(\mathbf{y}(t^\ell - 1)) \simeq \mathbf{c}^\ell$$

which is used for computing the one-step-ahead prediction at time t^ℓ . All the details remain the same when applying this method to equation (9) for iEEG data, except that $\mathbf{y}(t^\ell - 1)$ and $\mathbf{y}(t^m - 1)$ in equation (10), equation (11) and equation (12) are replaced with $[\mathbf{y}(t^\ell - 1)^T \dots \mathbf{y}(t^\ell - d)^T]^T$ and $[\mathbf{y}(t^m - 1)^T \dots \mathbf{y}(t^m - d)^T]^T$, respectively. We tuned h separately for fMRI and iEEG (see 'Hyper-parameter selection' below), giving rise to values that are so large that they essentially result in a globally linear model (**Supplementary Fig. 1**). The value of h was independently optimized for the computations reported in **Fig. 4i-j**, as described below.

Nonlinear minimum mean squared error models (optimal) ('MMSE (pairwise)', 'MMSE (scalar)'): The models in equation (5) (for fMRI) or equation (9) (for iEEG) ultimately define a stochastic mapping from $\mathbf{y}(t-1)$ or $(\mathbf{y}(t-1), \dots, \mathbf{y}(t-d))$ to $\mathbf{y}(t) - \mathbf{y}(t-1)$ such that observing the values of the former provides information to predict the latter. It is not hard to show that for two random

variables U and V , the optimal (i.e., minimum variance) prediction of U given $V = v$ is given by its conditional expectation $\hat{u} = E[U|V = v]$ known as the MMSE prediction.⁸⁰ Therefore, the optimal prediction of $\mathbf{y}(t) - \mathbf{y}(t-1)$ given $\mathbf{y}(t-1)$ or $(\mathbf{y}(t-1), \dots, \mathbf{y}(t-d))$ is given by $E[\mathbf{y}(t) - \mathbf{y}(t-1)|\mathbf{y}(t-1)]$ and $E[\mathbf{y}(t) - \mathbf{y}(t-1)|\mathbf{y}(t-1), \dots, \mathbf{y}(t-d)]$, respectively. Due to its optimality, it provides a theoretical upper bound on the achievable accuracy of *any* nonlinear model. The difficulty in calculating this estimate, however, is the estimation of the conditional distribution of $\mathbf{y}(t) - \mathbf{y}(t-1)$ given an observation of $\mathbf{y}(t-1)$ or $(\mathbf{y}(t-1), \dots, \mathbf{y}(t-d))$. Without imposing additional assumptions (e.g., linearity or Gaussianity), this task is not feasible in $n \sim 100$ dimensions with our limited data points per recording segment. However, this distribution is indeed feasible (i) on a pairwise basis, giving us the optimal predictions $E[y_i(t) - y_i(t-1)|y_j(t-1)]$ for all pairs $i, j = 1, \dots, n$, or (ii) on a scalar AR basis, yielding the optimal predictions $E[y_i(t) - y_i(t-1)|y_i(t-1), \dots, y_i(t-d)]$ separately for each $i = 1, \dots, n$. We use the former for fMRI and the latter for iEEG. To estimate this conditional distribution for fMRI, we use a Gaussian window with a standard deviation equal to β times the range of $y_j(t)$ in the training data to detect the training points close to each test $y_j(t-1)$ and then use an N -point weighted histogram to estimate the (conditional) distribution. More precisely, for any pair of i and j and any given *test* time t^ℓ , we first compute the Gaussian weight

$$w_j^{\ell, m} = \exp\left(-\frac{[y_j(t^m - 1) - y_j(t^\ell - 1)]^2}{2\sigma^2}\right)$$

$$\sigma = \beta \cdot \left(\max_{j, t^m} y_j(t^m) - \min_{j, t^m} y_j(t^m)\right)$$

for all *training* times t^m . Note that σ is obtained only using training data. Then, we divide the interval

$$\left[\min_{i, t^m} y_i(t^m) - y_i(t^m - 1), \max_{i, t^m} y_i(t^m) - y_i(t^m - 1)\right]$$

into N equal bins and construct a weighted histogram of all the *training* $y_i(t^m) - y_i(t^m - 1)$ where we count each $y_i(t^m) - y_i(t^m - 1)$ as much as $w_j^{\ell, m}$. We then normalize this histogram by dividing all bin values by their sum so that we obtain a well-defined (conditional) probability distribution $p_{i,j}^\ell(\Delta y_i^k)$, and estimate the expected value of this probability distribu-

tion as

$$E[y_i(t) - y_i(t-1)|y_j(t-1)] \simeq \sum_{k=1}^N \Delta y_i^k \cdot p_{i,j}^\ell(\Delta y_i^k)$$

where Δy_i^k denotes the center of the k 'th histogram bin. The case for iEEG is similar, except that the Gaussian weights are computed based on d -dimensional Euclidean distances given that we condition on d -dimensional vectors. In both cases, β and N are hyper-parameters that are tuned separately for fMRI and iEEG (see 'Hyper-parameter selection' below).

Zero model ('Zero'): So far, we have discussed several families of models. Comparisons among them will provide a clear picture of which family provides the best fit to the data *relative to* the others. Note though that this process does not necessarily imply that the best model is good in any absolute sense. In other words, all models may be estimating $\hat{\mathbf{y}}(t|t-1)$ at chance level or lower. Therefore, we also consider the zero model (a.k.a. zero-order hold, naive model, or random walk)

$$\mathbf{y}(t) - \mathbf{y}(t-1) = \mathbf{e}(t)$$

with the trivial estimate $\hat{\mathbf{y}}(t|t-1) = \mathbf{y}(t-1)$. Note that this expression corresponds to equation (5) with $f(\mathbf{y}(t-1)) = \mathbf{0}$ and is only meant to provide a baseline for comparison, not to act as a formal model itself. Also note that this model is different from, and often performs better than, the constant predictor $\hat{\mathbf{y}}(t|t-1) = \bar{\mathbf{y}}$ which constitutes the denominator of R^2 .

Hyper-parameter selection. For all models that involve the choice of a design hyper-parameter, we simultaneously optimized over all the hyper-parameters using stochastic gradient descent (SGD) with minibatch, separately for fMRI and iEEG. Let N_{param} denote the number of hyper-parameters in any of the models. Starting from an initial estimate of the hyper-parameter vector, in each iteration, $3^{N_{\text{param}}}$ hyper-parameter vectors were generated, constituting a hyper-cubic mesh around the current hyper-parameter estimate. For integer-valued hyper-parameters, we moved 1 point in each direction while for real-valued hyper-parameters, we moved 10^{-6} units. Using a minibatch of randomly selected data segments, the mean-over-minibatch of

the median-over-regions of the model R^2 was computed and maximized over the mesh. The random minibatch selection was independent between mesh points and between iterations. For integer-valued hyper-parameters, their value was updated to that of the maximizing mesh point. For real-value hyper-parameters, a gradient-ascent step was taken in the direction of the largest R^2 . The process was repeated until the hyper-parameters stopped having a consistent decrease/increase and hovered around a steady state value (which always happens due to the stochastic nature of SGD) and/or the R^2 stopped having a consistent increase.

For the ‘DNN (CNN)’ and ‘Subspace’ methods (the latter only in iEEG data), the aforementioned procedure was infeasibly slow. This was the case due to a high number of hyper-parameters for the CNN model and due to cluster jobs becoming frequently hung and needing to be killed and restarted for the subspace method. As such, we slightly modified the above procedure as follows. In each iteration of the SGD, instead of generating $3^{N_{\text{param}}}$ search directions at all the points of a hyper-cubic mesh, we generated $2N_{\text{param}} + 1$ search directions, one at the current optimum estimate and 2 at the current estimate $\pm 1\mathbf{e}_k$ (for integer-valued) or $\pm 10^{-6}\mathbf{e}_k$ (for real-valued) for each k ’th hyper-parameter, where \mathbf{e}_k is the k ’th canonical unit vector (all zeros except one 1 at the k ’th location). The remaining details were similar to the general case above. Finally, in all cases, we used 100 participants for fMRI (out of the total of 700) and 1500 segments for iEEG (out of the total of 8490) for hyper-parameter tuning and then removed them from the subsequent model fitting and validation experiments to ensure a lack of over-fitting to hyper-parameters.

The hyper-parameter and R^2 values throughout the process are shown in **Supplementary Figs. 7,8** for fMRI and iEEG data, respectively, and the final values of the hyper-parameters selected for each model are reported in **Supplementary Table 1**. Note that the initial hyper-parameter estimates were chosen based on prior experience, not randomly, which is why they are often very close to or the same as the final values.

Cross-validation. For the comparisons of **Fig. 2** on HCP data, we performed the cross-validation as follows, with slightly different procedures for brain-

wide and regional methods. For the brain-wide methods, for each of the 700 participants, we split each of the 4 resting scans of that participant into 2 halves, giving a total of 8 segments, each of length 600 samples. All of our methods were then applied using an 8-fold cross-validation where each time one of the 8 segments was used for testing and the remaining 7 were used for training. For pairwise methods, we were forced to lower the sample size due to the extremely high computational complexity of the MMSE predictor. Therefore, instead of each of the above 8 segments (per participant), we used the second quarter of that segment, giving us still an 8-fold cross-validation but on segments of length 150 samples each.

For the comparisons of **Fig. 3** on RAM data, we first split each of the 283, 5-minute recordings into 8490, 10-second segments. Even though having longer segments would in principle benefit model fitting, 10-second segments ensured that all of our methods could run using the 64 GB available memory per node on the CUBIC cluster. From the 8490 segments, those that contained any NaN entries (316 segments) or for which the subspace method produced NaN predictions (30 segments, due to the bad conditioning of the Φ matrix therein) were removed from further analysis. Since each recording is already split into 30 segments, and due to the large number of segments, we performed only a single-fold cross-validation on each segment, with the first 8 seconds used for training and the final 2 second used for cross-validation.

Multivariate test of whiteness. A standard measure of the goodness of fit in the prediction error method is the whiteness of residuals, measuring the extent to which all temporal structure (i.e., dynamics) in the data has been captured by the model. Note that a multivariate time series $\mathbf{e}(t)$ is “white” if it has no statistical dependence across time (i.e., $\mathbf{e}(s)$ and $\mathbf{e}(t)$ are independent if $s \neq t$) even though it can have arbitrary statistical dependence across channels (i.e., $e_i(t)$ and $e_j(t)$ can be dependent at the same time t). Parametric (χ^2) statistical tests have been devised for multivariate whiteness, such as the classical Box-Pierce portmanteau test⁸¹ and its modifications by Ljung & Box⁸² and Li & McLeod.⁸³ Under strong assumptions, all of these tests have a statistic Q (defined slightly differently between them) that

is asymptotically (at infinite samples) χ^2 distributed. In our datasets we found, however, that Q is not χ^2 distributed, and therefore we use randomization to generate the true null distribution of Q by shuffling the time indices of $\mathbf{e}(t)$ 100 times, computing Q for each of them, and computing the 95th percentile of the randomized Q values as the threshold Q_{thr} for significance. We use the original definition of Q (Ref.⁸¹),

$$Q = (N - M) \sum_{i=1}^M \text{tr} \left(\hat{\mathbf{R}}_{\mathbf{e}}(i)^T \hat{\mathbf{R}}_{\mathbf{e}}(0)^{-1} \hat{\mathbf{R}}_{\mathbf{e}}(i) \hat{\mathbf{R}}_{\mathbf{e}}(0)^{-1} \right)$$

where N is the number of (test) samples, M is the number of cross-correlation lags, and

$$\hat{\mathbf{R}}_{\mathbf{e}}(i) = \frac{1}{N - M} \sum_{t=0}^{N-M-1} \mathbf{e}(t+i) \mathbf{e}(t)^T, i = 0, 1, \dots, M.$$

is a finite-sample estimate of the cross-correlation matrix between channels of $\mathbf{e}(t)$ at lag i . Since in practice $\hat{\mathbf{R}}_{\mathbf{e}}(0)$ may be singular or near-singular, we use the pseudo-inverse of $\hat{\mathbf{R}}_{\mathbf{e}}(0)$ instead of its inverse in computing Q . Finally, only in the case of “pairwise” fMRI models where the residuals are inherently univariate, we use the simpler χ^2 test of whiteness for univariate time series (Ref.²⁵ Sec. 16.6).

Nonlinear predictors used for the analysis of the linearizing effects of macroscopic dynamics. Our discussions of linear and nonlinear models and their hyper-parameters so far applies to the comparisons shown in **Figs. 2,3** on neuroimaging time series. In our numerical analysis of the linearizing effects of macroscopic brain dynamics in **Fig. 4**, we also construct linear and nonlinear predictors and compute their R^2 . The linear predictor is always a simple linear regression model, while the nonlinear predictor is the MMSE predictor for two-dimensional predictions (**Fig. 4a-h**), and the manifold-based predictor for higher-dimensional predictions (**Fig. 4i-j**). The MMSE predictor was as above, except that β was adjusted as $0.02 + 0.02/\text{SNR}$ for **Fig. 4g-h**. For the manifold-based predictor, we used a Gaussian window and swept logarithmically over its hyper-parameter h from 0.1 to 10 in every iteration and chose the value of h that gave the largest R^2 . **Fig. 4j**-left panel shows the average of the resulting optimal h for 100 iterations.

Simulations involving the Izhikevic model. To show the linearizing effects spatiotemporal averaging

on a model whose ground truth is known to be nonlinear, we generated simulated time series from the Izhikevic model²³

$$\dot{v}(t) = 0.04v(t)^2 + 5v(t) + 140 - u(t) + I$$

$$\dot{u}(t) = a(bv(t) - u(t))$$

$$(v(t), u(t)) \leftarrow (c, u(t) + d) \quad \text{if } v(t) \geq 30$$

with $a = 0.02$, $b = 0.2$, $c = -65$, $d = 2$, $I = 7$. We discretize the model using Euler discretization with 0.1ms sampling.

Estimation of rsfMRI SNR. Here we describe our method for the estimation of rsfMRI time series scanner noise and the resulting SNR reported in **Supplementary Fig. 4**. From the 700 participants used for the study, 50 were selected uniformly at random, and for each selected participant, one of their 4 rest scans was selected also uniformly at random. The following was then performed for each of the 50 participant-scans. The rest scan was motion corrected using intra-modal linear registration with 6 degrees of freedom (in general, we kept the amount of pre-processing as minimal as possible throughout the SNR estimation algorithm since each preprocessing step often involves averaging and/or interpolation steps that can bias SNR estimates). The first volume of the motion-corrected rest scan was visually inspected and 10 voxels outside of the head were selected. Due to the unavailability of phantom scans, we used these voxels to estimate the scanner noise, while the two (phantom scans and outside voxels) have been shown to yield consistent noise estimates.⁸⁴ For each of the 10 voxels, we calculated the temporal variance of the corresponding time series and averaged the results, providing an estimate of scanner noise variance σ_N^2 . To estimate the signal power, a gray matter mask was extracted using each participant’s T1 scan and linearly registered back to the participant’s motion corrected rest scan. We then computed the temporal variance of each gray matter voxel and averaged the results, yielding an estimate of the combined signal and noise variance. Assuming statistical independence between scanner noise and the participants’ BOLD activity, this combined variance is precisely the sum $\sigma_S^2 + \sigma_N^2$ of signal variance and noise variance. The SNR was then calculated as σ_S/σ_N . Note that this process is inherently conservative and provides an upper bound on the SNR, as it,

for instance, does not include any physiological signals into “noise”. Therefore, the ratio between the power of signals of neural origin over all other signals contributing to rsfMRI time series may be much lower than 6.5. An SNR of about 6.5, however, is still low enough to yield a notable linearizing effect, highlighting the importance of measurement noise in downstream computational modeling.

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

The fMRI and iEEG data that support the findings of this study are publicly available, respectively, from the HCP S1200 Release at <https://www.humanconnectome.org/study/hcp-young-adult/document/1200-subjects-data-release>, and the RAM Public Data Release at <http://memory.psych.upenn.edu/RAM>.

Code availability

All the computations of this study (after data pre-processing) are performed in MATLAB, and the corresponding codes are publicly available at <https://github.com/enozari/rest-system-id>.

References

- [1] Kriegeskorte, N. & Douglas, P. K. Cognitive computational neuroscience. *Nat Neurosci* **21**, 1148–1160 (2018).
- [2] Wilson, R. C. & Niv, Y. Is model fitting necessary for model-based fMRI? *PLoS Comput Biol* **11**, e1004237 (2015).
- [3] Ruff, D. A., Ni, A. M. & Cohen, M. R. Cognition as a window into neuronal population space. *Annu Rev Neurosci* **41**, 77–97 (2018).
- [4] Hurwitz, C., Kudryashova, N., Onken, A. & Hennig, M. H. Building population models for large-scale neural recordings: Opportunities and pitfalls. *Current Opinion in Neurobiology* **70**, 64–73 (2021).
- [5] Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation through neural population dynamics. *Annual Review of Neuroscience* **43**, 249 (2020).
- [6] Amunts, K. *et al.* The human brain project: creating a european research infrastructure to decode the human brain. *Neuron* **92**, 574–581 (2016).
- [7] Gu, S. *et al.* Controllability of structural brain networks. *Nature communications* **6**, 1–10 (2015).
- [8] Sani, O. G. *et al.* Mood variations decoded from multi-site intracranial human brain activity. *Nature biotechnology* **36**, 954–961 (2018).
- [9] Izhikevich, E. M. *Dynamical Systems in Neuroscience* Computational neuroscience Dynamical systems in neuroscience (MIT Press, 2007). URL <https://books.google.com/books?id=kVjM6DFk-twC>.
- [10] Booth, V. & Rinzel, J. A minimal, compartmental model for a dendritic origin of bistability of motoneuron firing patterns. *J Comput Neurosci* **2**, 299–312 (1995).
- [11] Freeman, W. J. Nonlinear gain mediating cortical stimulus-response relations. *Biological Cybernetics* **33**, 237–247 (1979).
- [12] Wilson, H. R. & Cowan, J. D. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal* **12**, 1–24 (1972).
- [13] Li, X., Coyle, D., Maguire, L., McGinnity, T. M. & Benali, H. A model selection method for nonlinear system identification based fmri effective connectivity analysis. *IEEE transactions on medical imaging* **30**, 1365–1380 (2011).
- [14] Wang, Y. M., Schultz, R. T., Constable, R. T. & Staib, L. H. *Nonlinear estimation and modeling of fmri data using spatio-temporal support vector regression*, 647–659 (Springer, 2003).
- [15] Stephan, K. E. *et al.* Nonlinear dynamic causal models for fmri. *Neuroimage* **42**, 649–662 (2008).

- [16] Ritter, P., Schirner, M., McIntosh, A. R. & Jirsa, V. K. The virtual brain integrates computational modeling and multimodal neuroimaging. *Brain Connect* **3**, 121–145 (2013).
- [17] Stam, C. J. Nonlinear dynamical analysis of eeg and meg: review of an emerging field. *Clinical neurophysiology* **116**, 2266–2301 (2005).
- [18] Ehlers, C. L., Havstad, J., Prichard, D. & Theiler, J. Low doses of ethanol reduce evidence for nonlinear structure in brain activity. *Journal of Neuroscience* **18**, 7474–7486 (1998).
- [19] Gultepe, E. & He, B. A linear/nonlinear characterization of resting state brain networks in fmri time series. *Brain topography* **26**, 39–49 (2013).
- [20] Blinowska, K. J. & Malinowski, M. Non-linear and linear forecasting of the eeg time series. *Biological cybernetics* **66**, 159–165 (1991).
- [21] Zhao, Y., Billings, S. A., Wei, H.-L. & Sarri-
giannis, P. G. A parametric method to measure time-varying linear and nonlinear causality with applications to eeg data. *IEEE Transactions on Biomedical Engineering* **60**, 3141–3148 (2013).
- [22] Yang, Y., Sani, O. G., Chang, E. F. & Shanechi, M. M. Dynamic network modeling and dimensionality reduction for human ecog activity. *Journal of neural engineering* **16**, 056014 (2019).
- [23] Izhikevich, E. M. Simple model of spiking neurons. *IEEE Transactions on neural networks* **14**, 1569–1572 (2003).
- [24] Gorrostieta, C., Fiecas, M., Ombao, H., Burke, E. & Cramer, S. Hierarchical vector autoregressive models and their applications to multi-subject effective connectivity. *Frontiers in computational neuroscience* **7**, 159 (2013).
- [25] Ljung, L. System identification: theory for the user. *PTR Prentice Hall, Upper Saddle River, NJ* 1–14 (1999).
- [26] Kim, S. A. & Ching, S. *Quasilinearization-based controllability analysis of neuronal rate networks*, 7371–7376 (IEEE, 2016).
- [27] Buzsáki, G., Anastassiou, C. A. & Koch, C. The origin of extracellular fields and currents—eeg, ecog, lfp and spikes. *Nature reviews neuroscience* **13**, 407–420 (2012).
- [28] Lindén, H., Pettersen, K. H. & Einevoll, G. T. Intrinsic dendritic filtering gives low-pass power spectra of local field potentials. *Journal of computational neuroscience* **29**, 423–444 (2010).
- [29] Greve, D. N., Brown, G. G., Mueller, B. A., Glover, G. & Liu, T. T. A survey of the sources of noise in fmri. *Psychometrika* **78**, 396–416 (2013).
- [30] Liu, Y., Coon, W., De Pestors, A., Brunner, P. & Schalk, G. The effects of spatial filtering and artifacts on electrocorticographic signals. *Journal of neural engineering* **12**, 056008 (2015).
- [31] Yang, Z. *Incorporating Structural Bias into Neural Networks for Natural Language Processing*. Ph.D. thesis, Carnegie Mellon University (2019).
- [32] Kononova, A. V., Corne, D. W., De Wilde, P., Shneer, V. & Caraffini, F. Structural bias in population-based algorithms. *Information Sciences* **298**, 468–490 (2015).
- [33] Mehta, R., Shen, C., Xu, T. & Vogelstein, J. T. A consistent independence test for multivariate time-series. *arXiv preprint arXiv:1908.06486* (2019).
- [34] Dafilis, M. P., Sinclair, N. C., Cadusch, P. J. & Liley, D. T. Re-evaluating the performance of the nonlinear prediction error for the detection of deterministic dynamics. *Physica D: Nonlinear Phenomena* **240**, 695–700 (2011).
- [35] Deneux, T. & Faugeras, O. Using nonlinear models in fmri data analysis: model selection and activation detection. *NeuroImage* **32**, 1669–1689 (2006).
- [36] Liu, Z. *et al.* Linear and nonlinear relationships between visual stimuli, eeg and bold fmri signals. *Neuroimage* **50**, 1054–1066 (2010).
- [37] Schulz, M.-A. *et al.* Different scaling of linear models and deep learning in ukbiobank brain images versus machine-learning datasets. *Nature communications* **11**, 1–15 (2020).
- [38] He, T. *et al.* Deep neural networks and kernel regression achieve comparable accuracies for functional connectivity prediction of behavior and demographics. *NeuroImage* **206**, 116276 (2020).

- [39] Wobst, P., Wenzel, R., Kohl, M., Obrig, H. & Villringer, A. Linear aspects of changes in deoxygenated hemoglobin concentration and cytochrome oxidase oxidation during brain activation. *Neuroimage* **13**, 520–530 (2001).
- [40] Rumyantsev, O. I. *et al.* Fundamental bounds on the fidelity of sensory cortical coding. *Nature* **580**, 100–105 (2020).
- [41] Khalil, H. K. *Nonlinear Systems* Pearson Education (Prentice Hall, 2002). URL https://books.google.com/books?id=t_d1QgAACAAJ.
- [42] Palva, J. M. & Palva, S. Functional integration across oscillation frequencies by cross-frequency phase synchronization. *Eur J Neurosci* **48**, 2399–2406 (2018).
- [43] Zañudo, J. G. T., Yang, G. & Albert, R. Structure-based control of complex networks with nonlinear dynamics. *Proc Natl Acad Sci U S A* **114**, 7234–7239 (2017).
- [44] Rozum, J. C. & Albert, R. Identifying (un)controllable dynamical behavior in complex networks. *PLoS Comput Bio* **14**, e1006630 (2018).
- [45] Tang, E. & Bassett, D. S. Colloquium: Control of dynamics in brain networks. *Reviews of modern physics* **90**, 031003 (2018).
- [46] Towlson, E. K. *et al.* Caenorhabditis elegans and the network control framework-FAQs. *Philos Trans R Soc Lond B Biol Sci* **373**, 20170372 (2018).
- [47] Karrer, T. M. *et al.* A practical guide to methodological considerations in the controllability of structural brain networks. *J Neural Eng* **17**, 026031 (2020).
- [48] Goutte, C., Nielsen, F. A. & Hansen, L. K. Modeling the hemodynamic response in fmri using smooth fir filters. *IEEE transactions on medical imaging* **19**, 1188–1201 (2000).
- [49] Haller, M. *et al.* Parameterizing neural power spectra. *BioRxiv* 299859 (2018).
- [50] Bansal, K., Nakuci, J. & Muldoon, S. F. Personalized brain network models for assessing structure-function relationships. *Curr Opin Neurobiol* **52**, 42–47 (2018).
- [51] Schirner, M., Rothmeier, S., Jirsa, V. K., McIntosh, A. R. & Ritter, P. An automated pipeline for constructing personalized virtual brains from multimodal neuroimaging data. *Neuroimage* **117**, 343–357 (2015).
- [52] Bayrak, S., Hövel, P. & Vuksanović, V. in *Modeling functional connectivity on empirical and randomized structural brain networks* (Springer, 2017).
- [53] Saenger, V. M. *et al.* Uncovering the underlying mechanisms and whole-brain dynamics of deep brain stimulation for parkinson’s disease. *Scientific reports* **7**, 1–14 (2017).
- [54] Zarghami, T. S. & Friston, K. J. Dynamic effective connectivity. *Neuroimage* **207**, 116453 (2020).
- [55] Friston, K. J. *et al.* DCM for complex-valued data: cross-spectra, coherence and phase-delays. *Neuroimage* **59**, 439–455 (2012).
- [56] Becker, C. O., Bassett, D. S. & Preciado, V. M. Large-scale dynamic modeling of task-fmri signals via subspace system identification. *Journal of neural engineering* **15**, 066016 (2018).
- [57] Yang, Y., Connolly, A. T. & Shanechi, M. M. A control-theoretic system identification framework and a real-time closed-loop clinical simulation testbed for electrical brain stimulation. *J Neural Eng* **15**, 066007 (2018).
- [58] Barch, D. M. Resting-state functional connectivity in the human connectome project: Current status and relevance to understanding psychopathology. *Harv Rev Psychiatry* **25**, 209–217 (2017).
- [59] Burgess, G. C. *et al.* Evaluation of denoising strategies to address motion-correlated artifacts in resting-state functional magnetic resonance imaging data from the human connectome project. *Brain Connect* **6**, 669–680 (2016).
- [60] Elam, J. Hcp data release updates: Known issues and planned fixes. <https://wiki.humanconnectome.org/display/PublicData/HCP+Data+Release+Updates%3A+Known+Issues+and+Planned+fixes> (2020). Last accessed May 13, 2020.

- [61] Schaefer, A. *et al.* Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebral Cortex* **28**, 3095–3114 (2018).
- [62] Tian, Y., Margulies, D. S., Breakspear, M. & Zalesky, A. Hierarchical organization of the human subcortex unveiled with functional connectivity gradients. *bioRxiv* (2020).
- [63] Stiso, J. *et al.* White matter network architecture guides direct electrical stimulation through optimal state transitions. *Cell Rep* **28**, 2554–2566.e7 (2019).
- [64] Khambhati, A. N. *et al.* Functional control of electrophysiological network architecture using direct neurostimulation in humans. *Netw Neurosci* **3**, 848–877 (2019).
- [65] Betzel, R. F. *et al.* Structural, geometric and genetic factors predict interregional brain connectivity patterns probed by electrocorticography. *Nat Biomed Eng* **3**, 902–916 (2019).
- [66] Ung, H. *et al.* Interictal epileptiform activity outside the seizure onset zone impacts cognition. *Brain* **140**, 2157–2168 (2017).
- [67] Ren, S., Gliske, S. V., Brang, D. & Stacey, W. C. Redaction of false high frequency oscillations due to muscle artifact improves specificity to epileptic tissue. *Clinical Neurophysiology* **130**, 976–985 (2019).
- [68] Owen, L. L. *et al.* A gaussian process model of human electrocorticographic data. *Cerebral Cortex* **30**, 5333–5345 (2020).
- [69] Prime, D., Rowlands, D., O’Keefe, S. & Dionisio, S. Considerations in performing and analyzing the responses of cortico-cortical evoked potentials in stereo-eeg. *Epilepsia* **59**, 16–26 (2018).
- [70] Delorme, A., Sejnowski, T. & Makeig, S. Enhanced detection of artifacts in eeg data using higher-order statistics and independent component analysis. *Neuroimage* **34**, 1443–1449 (2007).
- [71] Stiso, J. *et al.* Fluctuations in functional connectivity associated with interictal epileptiform discharges (ieds) in intracranial eeg. *bioRxiv* 2021–05 (2022).
- [72] Lawhern, V., Hairston, W. D., McDowell, K., Westerfield, M. & Robbins, K. Detection and classification of subject-generated artifacts in eeg signals using autoregressive models. *Journal of neuroscience methods* **208**, 181–189 (2012).
- [73] Bressler, S. L., Richter, C. G., Chen, Y. & Ding, M. Cortical functional network organization from autoregressive modeling of local field potential oscillations. *Statistics in medicine* **26**, 3875–3885 (2007).
- [74] Deshpande, R., Wu, G.-R., Marinazzo, D., Hu, X. & Deshpande, G. Hemodynamic response function (hrf) variability confounds resting-state fmri functional connectivity. *Magnetic resonance in medicine* **80**, 1697–1713 (2018).
- [75] Taylor, A. J., Kim, J. H. & Ress, D. Characterization of the hemodynamic response function across the majority of human cerebral cortex. *NeuroImage* **173**, 322–331 (2018).
- [76] Singh, M., Braver, T., Cole, M. & Ching, S. Individualized dynamic brain models: Estimation and validation with resting-state fmri. *bioRxiv* 678243 (2019).
- [77] Roll, J. *Local and Piecewise Affine Approaches to System Identification*. Ph.D. thesis, Linköping University (2003).
- [78] Ljung, L. *Approaches to identification of nonlinear systems*, 1–5 (IEEE, 2010).
- [79] Popivanov, D., Dushanova, J., Mineva, A. & Krekule, I. *Detection of successive changes in dynamics of eeg time series: linear and nonlinear approach*, Vol. 4, 1590–1591 (IEEE, 1996).
- [80] Poor, H. V. *An introduction to signal detection and estimation* (Springer Science & Business Media, 2013).
- [81] Box, G. E. & Pierce, D. A. Distribution of residual autocorrelations in autoregressive-integrated moving average time series models. *Journal of the American statistical Association* **65**, 1509–1526 (1970).
- [82] Ljung, G. M. & Box, G. E. P. On a measure of lack of fit in time series models. *Biometrika* **65**, 297–303 (1978).

- [83] Li, W. K. & McLeod, A. I. Distribution of the residual autocorrelations in multivariate arma time series models. *Journal of the Royal Statistical Society: Series B (Methodological)* **43**, 231–239 (1981).
- [84] Chen, C.-C. & Tyler, C. W. in *Spectral analysis of fmri signal and noise* 63–76 (Springer, 2008).
- [85] Mitchell, S. M., Lange, S. & Brus, H. Gendered citation patterns in international relations journals. *International Studies Perspectives* **14**, 485–492 (2013).
- [86] Dion, M. L., Sumner, J. L. & Mitchell, S. M. Gendered citation patterns across political science and social science methodology fields. *Political Analysis* **26**, 312–327 (2018).
- [87] Caplar, N., Tacchella, S. & Birrer, S. Quantitative evaluation of gender bias in astronomical publications from citation counts. *Nature Astronomy* **1**, 0141 (2017).
- [88] Maliniak, D., Powers, R. & Walter, B. F. The gender citation gap in international relations. *International Organization* **67**, 889–922 (2013).
- [89] Dworkin, J. D. *et al.* The extent and drivers of gender imbalance in neuroscience reference lists. *bioRxiv* (2020). URL <https://www.biorxiv.org/content/early/2020/01/11/2020.01.03.894378>.
- [90] Zhou, D. *et al.* Gender diversity statement and code notebook v1.0 (2020). URL <https://doi.org/10.5281/zenodo.3672110>.
- [91] Ambekar, A., Ward, C., Mohammed, J., Male, S. & Skiena, S. *Name-ethnicity classification from open sources*, 49–58 (2009).
- [92] Sood, G. & Laohaprapanon, S. Predicting race and ethnicity from the sequence of characters in a name. *arXiv preprint arXiv:1805.02109* (2018).

Acknowledgements

Funding. E.N. and D.S.B. disclose support for the research described in this study from the Swartz

Foundation and from the National Science Foundation (award number 1926757). J.S. discloses support for the publication of this study from the National Institute of Mental Health (award number 1F31MH120925-01A1). E.J.C. acknowledges support for the publication of this study from the National Institute of Mental Health (award number F30MH118871-01). D.S.B. also acknowledges support for the publication of this study from the John D. and Catherine T. MacArthur Foundation, the Institute for Scientific Interchange Foundation, the Paul G. Allen Family Foundation, and the Army Research Office (grant number W911NF-16-1-0474). The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

Author Contributions

E.N., G.J.P., and D.S.B. designed the research; E.N. performed the research; G.J.P., and D.S.B. supervised the research; M.A.B. preprocessed the fMRI data; J.S. preprocessed the iEEG data; M.A.B., J.S., L.C., E.J.C., X.H., and A.M. assisted in the neuroimaging analysis of the data; all authors contributed in writing the paper.

Competing interests

The authors declare no competing interests.

Citation Diversity Statement

Recent work in several fields of science has identified a bias in citation practices such that papers from women and other minority scholars are under-cited relative to the number of such papers in the field.^{85–89} Here we sought to proactively consider choosing references that reflect the diversity of the field in thought, form of contribution, gender, race, ethnicity, and other factors. First, we obtained the predicted gender of the first and last author of each reference by using databases that store the probability of a first name being carried by a woman.^{89,90} By this measure (and excluding self-citations to the first and last authors of our current paper), our references contain 8.73% woman(first)/woman(last), 18.87% man/-

woman, 18.34% woman/man, and 54.06% man/-man. This method is limited in that a) names, pronouns, and social media profiles used to construct the databases may not, in every case, be indicative of gender identity and b) it cannot account for intersex, non-binary, or transgender people. Second, we obtained predicted racial/ethnic category of the first and last author of each reference by databases that store the probability of a first and last name being carried by an author of color.^{91,92} By this measure (and excluding self-citations), our references contain 14.39% author of color (first)/author of color(last), 15.47% white author/author of color, 23.76% author of color/white author, and 46.37% white author/white author. This method is limited in that a) names and Florida Voter Data to make the predictions may not be indicative of racial/ethnic identity, and b) it cannot account for Indigenous and mixed-race authors, or those who may face differential biases due to the ambiguous racialization or ethnicization of their names. We look forward to future work that could help us to better understand how to support equitable practices in science.

a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

©The Author(s), under exclusive licence to Springer Nature Limited 2023

Additional Information

Extended data is available for this paper at <https://doi.org/10.1038/s41551-02X-XXXX-X>.

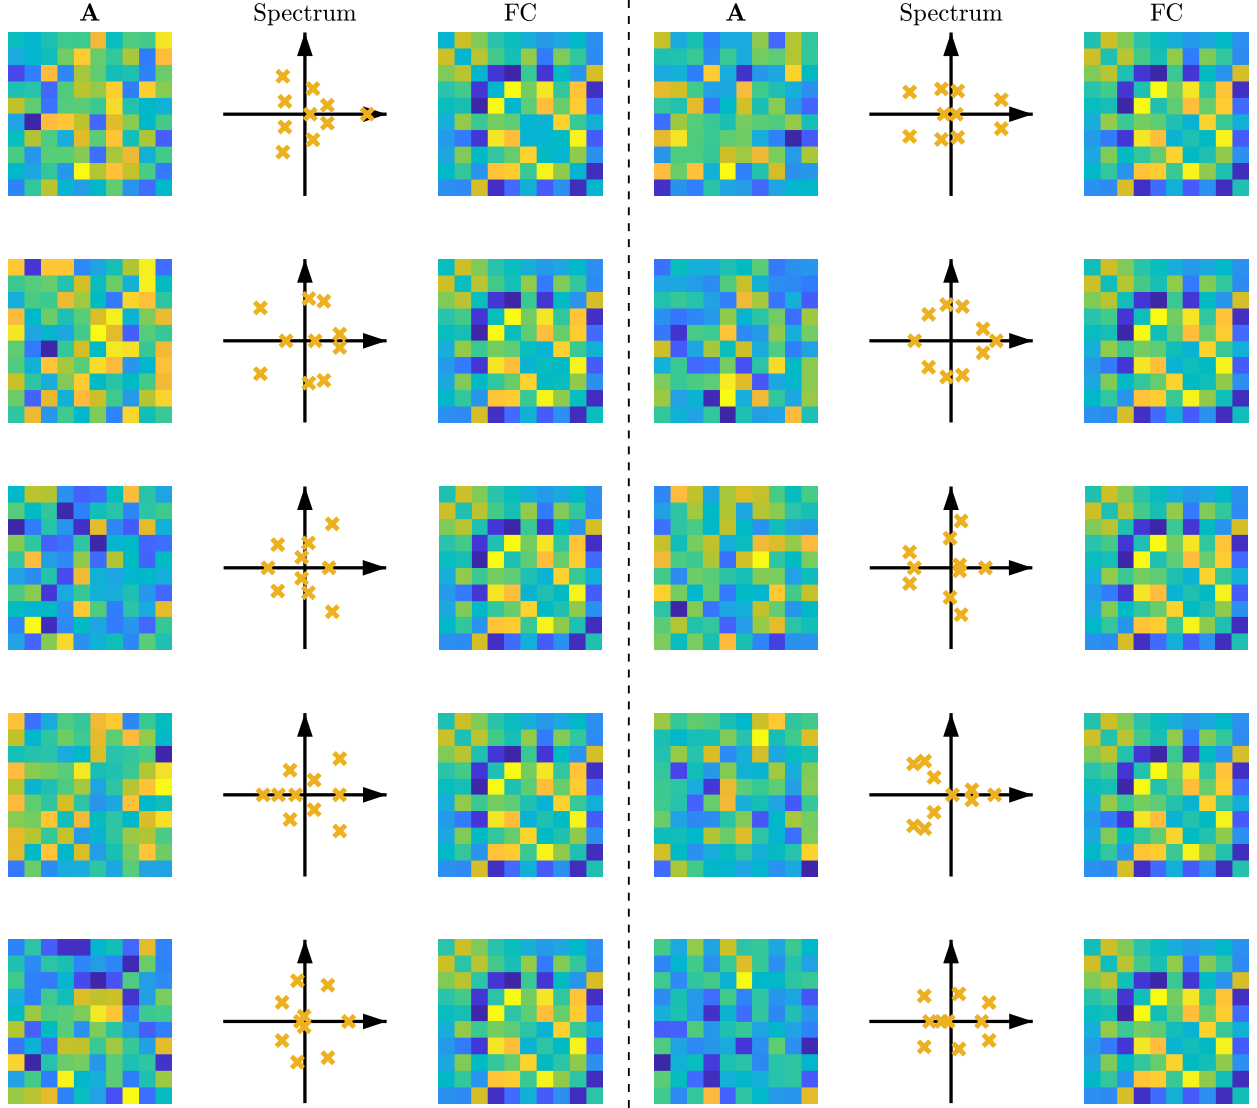
Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41551-02X-XXXX-X>.

Correspondence and requests for materials should be addressed to

Peer review information Nature Biomedical Engineering thanks Reviewer and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

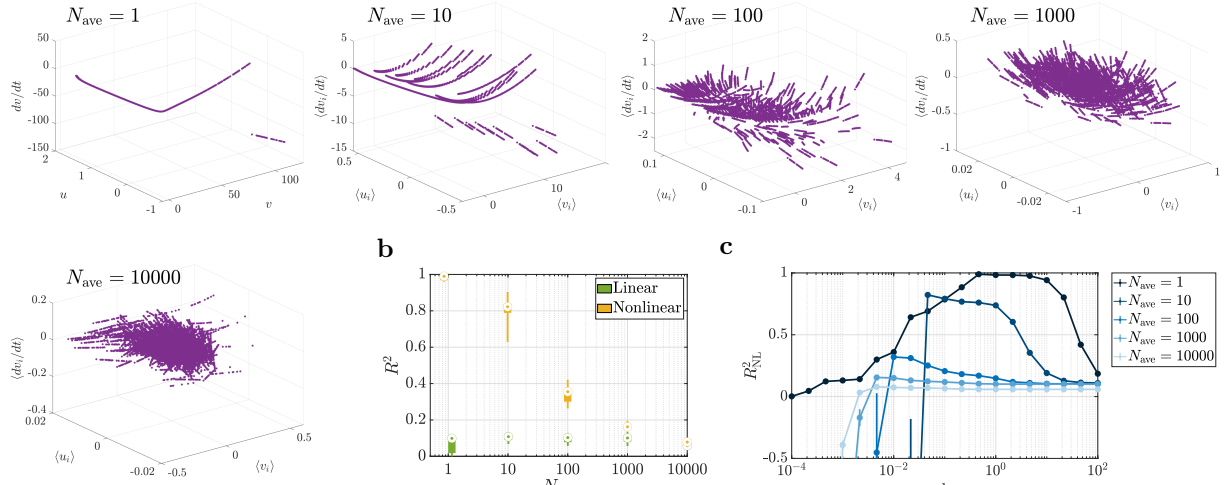
Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations. Springer Nature or its licensor holds exclusive rights to this article under

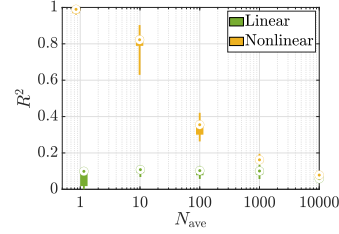


Extended Data Fig. 1 | Different systems can produce almost identical functional connectivity (FC). Ten different systems of the form $\mathbf{y}(t) = \mathbf{A}\mathbf{y}(t-1) + \mathbf{e}(t)$ are shown with completely different \mathbf{A} matrices and spectra (eigenvalues of \mathbf{A}), but almost identical FC matrices. The Pearson correlation coefficient between all pairs of FCs is 1.0000 while the absolute value of the Pearson correlation coefficients between the \mathbf{A} matrices range from a minimum of 0.0023 to a maximum of 0.36. All systems were driven by independent and randomly generated noise sequences $\mathbf{e}(t)$. In the figures, the diagonal entries of the FC matrices (which are by definition equal to 1) are set 0 to increase color resolution.

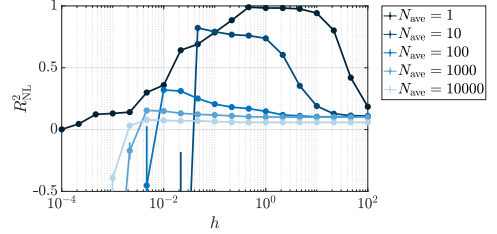
a



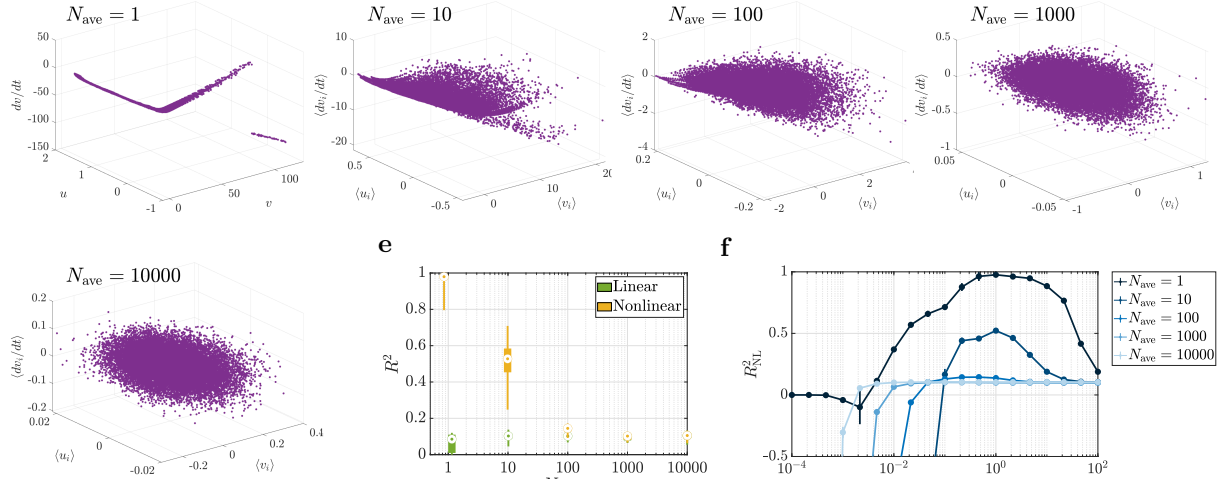
b



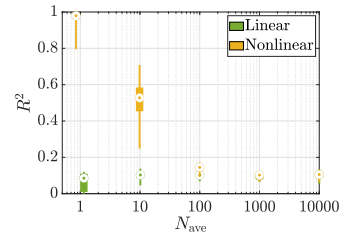
c



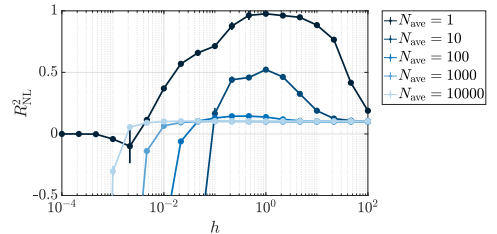
d



e



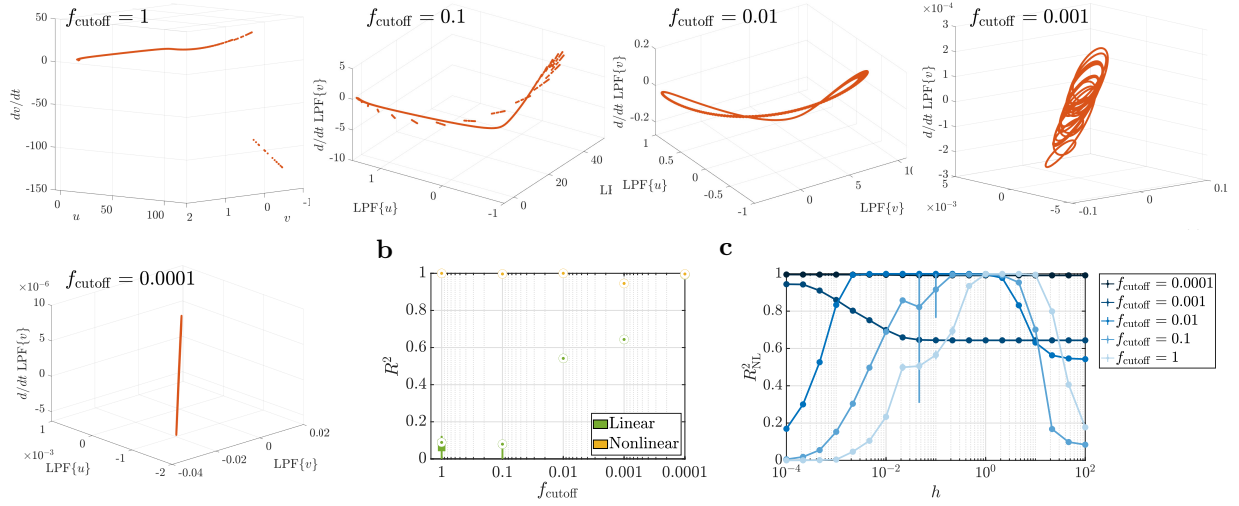
f



Extended Data Fig. 2 | The linearizing effect of spatial averaging on the Izhikevic spiking model.

(a) The Izhikevic model²³ was simulated for N_{ave} neurons and their v_i and u_i variables were averaged before being “observed”. The neurons had the same parameters (and hence the same spiking frequency) but were initialized at random phases of their limit cycles. The average of dv_i/dt is plotted against the average of v_i and u_i , showing the loss of nonlinearity at about $N_{\text{ave}} \sim 10^3$ - 10^4 . dv_i/dt is chosen since it has a more nonlinear dependence on v_i and u_i than du_i/dt for one neuron, but the same (even stronger) linearizing effect holds for $(\langle v_i \rangle, \langle u_i \rangle, \langle du_i/dt \rangle)$. (b) The R^2 distributions of a linear (simple linear regression) and nonlinear (manifold-based locally linear regression) model for the relationships in panel (a) and 100 random repetitions. The manifold-based model was chosen over the MMSE model as it consistently gave significantly higher R^2 values (due to the high sample complexity of the MMSE model). Note that the R^2 distribution of the linear model is mostly flat (less affected by spatial averaging), while that of the nonlinear model decays rapidly with N_{ave} until it reaches the linear level. (c) Nonlinear R^2 as a function of the manifold method’s window sizes h for varying N_{ave} values. The circles and error bars show median and inter-quartile ranges, respectively. While an optimal, mid-range h exists for small values of N_{ave} , the curve plateaus for large h as N_{ave} increases, showing the loss of nonlinearity with increasing N_{ave} . (d-f) Same as panels (a-c) but for a model where *process* noise is added to the dv_i/dt equations. The noise has a very low power (SNR = 100) and is only meant to destroy the artificial micro-nonlinear relationships that are visible in panel (a) and can inflate the power of nonlinear regression for $N_{\text{ave}} \sim 10$ - 10^3 . Note that this is different from the addition of observation noise discussed in the main text, even though adding observation noise could serve this purpose as well. In all box plots, the center point, box limits, and whiskers represent the median, upper and lower quartiles, and the smallest and largest samples, respectively.

a

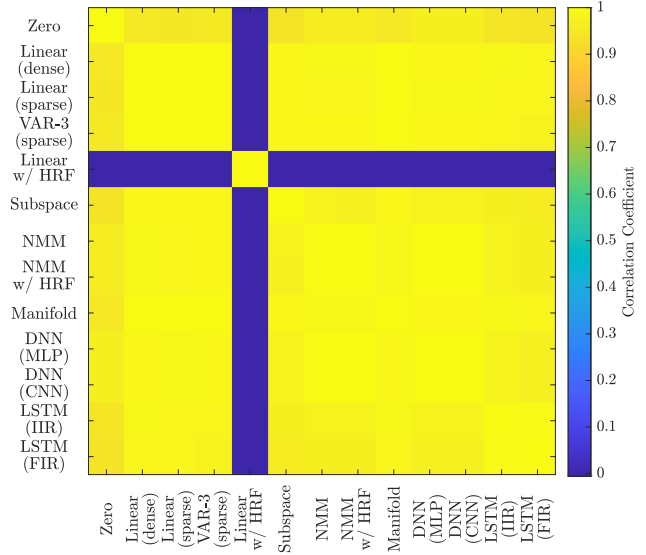
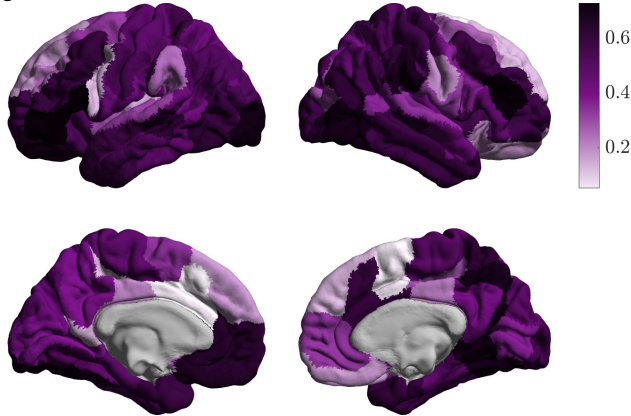
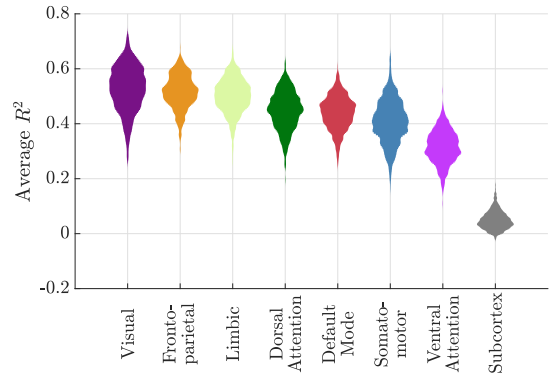


Extended Data Fig. 3 | The linearizing effect of temporal averaging on the Izhikevic spiking model.

Panels parallel those in Figure 2. **(a)** One Izhikevic model²³ was simulated, starting from a random initial condition on its limit cycle, and its v and u variables were low-pass filtered before being “observed”. $d/dt \text{ LPF}\{v\}$ is plotted against $\text{LPF}\{v\}$ and $\text{LPF}\{u\}$, showing the loss of nonlinearity at about $f_{\text{cutoff}} \sim 10^{-4}$. **(b,c)** Similar to panels (b,c) in Figure 2, except that here the R^2 of the linear model reaches that of the nonlinear one as temporal averaging is intensified. The nonlinear model always has a near-perfect prediction power given the fully deterministic nature of this simulation (i.e., a sufficiently small h was always sufficient to perfectly match a locally-linear model to the curves in panel (a)). Note that addition of process noise here does not have the same “blurring” effect of Figure 2(d-f) since it lies before the low-pass filter. Observation noise can nevertheless be added with the same effect (not shown here). In all box plots, the center point, box limits, and whiskers represent the median, upper and lower quartiles, and the smallest and largest samples, respectively.

a

Zero	0	1	1	0	0	0	0	6	3	8	9	1	59	61	65	50
Linear (dense)	99	0	69	42	51	5	10	98	81	100	100	10	100	100	100	100
Linear (sparse)	98	28	0	3	2	14	7	100	91	100	100	28	100	100	100	100
VAR-2 (sparse)	99	48	94	0	93	21	11	100	97	100	100	48	100	100	100	100
AR-2 (sparse)	99	42	96	6	0	18	7	100	96	100	100	42	100	100	100	100
VAR-3 (sparse)	99	93	85	78	79	0	71	98	97	100	100	93	100	100	100	100
AR-3 (sparse)	99	87	93	84	92	26	0	100	98	100	100	87	100	100	100	100
Linear w/ HRF	89	1	0	0	0	2	0	0	7	47	62	1	96	97	100	98
Subspace	92	11	4	2	3	3	0	93	0	91	94	11	98	98	98	98
NMM	89	0	0	0	0	0	0	46	5	0	51	0	91	93	97	94
NMM w/ HRF	86	0	0	0	0	0	0	31	3	45	0	0	91	91	93	91
Manifold	99	66	69	42	52	5	10	98	81	100	100	0	100	100	100	100
DNN (MLP)	16	0	0	0	0	0	0	3	2	6	9	0	0	63	75	59
DNN (CNN)	16	0	0	0	0	0	0	1	2	4	9	0	30	0	76	34
LSTM (IIR)	25	0	0	0	0	0	0	0	1	2	3	0	22	22	0	0
LSTM (FIR)	35	0	0	0	0	0	0	0	2	3	9	0	39	52	100	0
Zero																
Linear (dense)																
Linear (sparse)																
VAR-2 (sparse)																
AR-2 (sparse)																
VAR-3 (sparse)																
AR-3 (sparse)																
Linear w/ HRF																
Subspace																
NMM																
NMM w/ HRF																
Manifold																
DNN (MLP)																
DNN (CNN)																
LSTM (IIR)																
LSTM (FIR)																

b**c****d**

Extended Data Fig. 4 | The spatial distribution of predictability in fMRI data. (a) Regionwise comparisons between all pairs of models. The color and number in each cell indicates the percentage of brain regions where the median of the R^2 values of the model on the row is statistically significantly greater than the median of the R^2 values of the model on the column ($\alpha = 0.05$, Wilcoxon signed rank test, and Bonferroni correction for multiple comparisons). (b) The correlation coefficient between the R^2 values of all the (brain-wise) methods. All methods produce almost the same cortical distributions of R^2 (albeit with different absolute values of R^2 , cf. Figure 2a in the main text), except for 'Linear w/ HRF' which has an un-correlated R^2 distribution relative to the rest of the methods. (c) The cortical distribution of the R^2 of our best model ('VAR-3 (sparse)'), averaged over the 700 subjects. (d) Violin plots of the distribution of R^2 , averaged over all the regions of each resting state network, for all subjects. Note that each distribution in panel (b) is thus composed of 700 samples. All pairs of distributions have significantly different medians in the order plotted (one-sided Mann-Whitney U -test at $\alpha = 0.05$ with BH-FDR correction for multiple comparisons). The most striking difference is between the cortical and subcortical regions, where the dynamics of the latter are remarkably less predictable than the former. This lower predictability can be viewed as having "more noisy" dynamics or, more precisely, less spatially and temporally correlated (i.e., more white) fMRI time series in these regions.

Contents

1	Supplementary Note 1: Effect of Bandpass Filtering on System Identification and the Detection of Nonlinearity	2
2	Supplementary Table 1: List of linear and nonlinear families of models	4
3	Supplementary Figures	5

List of Supplementary Figures

1	Effect of window size h on the accuracy of the manifold-based locally linear ('Manifold') model in fMRI data.	5
2	The effects of the number of lags and sparsity patterns on the prediction accuracy and computational complexity of linear AR models of rsfMRI.	6
3	The effect of the LASSO parameter λ on the accuracy of the 'VAR-3 (sparse)' model.	7
4	Histogram of scanner SNR estimates for rsfMRI data.	8
5	Separate training and test times for system identification methods of rsfMRI. . .	9
6	Separate training and test times for system identification methods of rsiEEG. . .	10
7	Hyper-parameter tuning of linear and nonlinear model families for fMRI.	12
8	Hyper-parameter tuning of linear and nonlinear model families for iEEG.	14
9	Comparing the 'DNN (MLP)' model with linear models in fitting simulated data from the Izhikevic model.	15
10	Comparison of 'Linear (dense)' and 'DNN (MLP)' models on the logistic map dynamical system.	16
11	Comparing 'DNN (MLP)' model with a linear activation function against 'Linear (dense)' and 'Linear (sparse)'.	17
12	Linear vs. nonlinear models of finely-parcellated rsfMRI activity.	18
13	Linear vs. nonlinear models of unparcellated cortical rsfMRI activity.	20
14	Linear vs. nonlinear models of unparcellated subcortical rsfMRI activity.	21
15	Linear vs. nonlinear models of minimally pre-processed rsfMRI activity.	22
16	Test-retest validation of model comparisons for rsfMRI data.	23
17	The channel-wise R^2 distribution of the zero model for iEEG data with different subsampling ratios and the corresponding sampling frequency.	24
18	Linear versus nonlinear models of 5-fold subsampled rsiEEG activity.	25
19	Linear versus nonlinear models of 25-fold subsampled rsiEEG activity.	26
20	k -step ahead prediction of rsfMRI data.	27
21	k -step ahead prediction of rsiEEG data.	28
22	Model comparisons on iEEG data without channel removal.	29

1 Supplementary Note 1: Effect of Bandpass Filtering on System Identification and the Detection of Nonlinearity

A standard step in the preprocessing of resting state fMRI time series is bandpass filtering, typically over the range $[0.01, 0.08]$ Hz [1], to reduce the contribution of non-neuronal sources on the signal and improve the SNR. In this work, however, we purposefully avoided this step. In the following, we discuss in detail the role and effects of pre-filtering in this rather unconventional context of system identification and, in particular, detection of nonlinear dynamics.

1. First, band-pass filtering (or any linear filtering for this matter) has no effect on the fitting or evaluation of linear models. The reason, in short, is the commuting property of linear systems. More specifically, any linear system including all of the ones used in this work can be written in the impulse-response form [2]

$$\mathbf{y}(t) = \mathbf{G}(q)\mathbf{e}(t),$$

where \mathbf{G} is the transfer matrix

$$\mathbf{G}(q) = \begin{bmatrix} G_{11}(q) & G_{12}(q) & \cdots \\ G_{21}(q) & G_{22}(q) & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

such that for any i and j

$$G_{ij}(q) = g_{ij}(0) + g_{ij}(1)q^{-1} + g_{ij}(2)q^{-2} + \cdots$$

Here, $g_{ij}(t)$ is the impulse response from the the j th input $e_j(t)$ to the i th output $y_i(t)$, and q is the standard shift operator such that $q^{-1}s(t) = s(t-1)$ for any signal $s(t)$ [3]. Recall that $\mathbf{y}(t)$ is the BOLD time series without band-pass filtering, as used in the main text, and let

$$F(q) = f(0) + f(1)q^{-1} + f(2)q^{-2} + \cdots$$

be any linear filter, including the bandpass filter used in common preprocessing pipelines. Assume, without loss of generality, that $f(0) \neq 0$ (an almost identical argument can be given if $f(0)$ or any number of the first terms in $F(q)$ are zero, simply by factoring out enough powers of q^{-1}). The output of the filter is

$$\mathbf{z}(t) = F(q)\mathbf{y}(t) = F(q)\mathbf{G}(q)\mathbf{e}(t).$$

It then immediately follows from the prediction error framework [3] that the one step ahead prediction error of $\mathbf{z}(t)$ is given by

$$\begin{aligned} \mathbf{z}(t) - \hat{\mathbf{z}}(t|t-1) &= f(0)\mathbf{G}^{-1}(q)F^{-1}(q)\mathbf{z}(t) \\ &= f(0)\mathbf{G}^{-1}(q)\mathbf{y}(t) \\ &= f(0)[\mathbf{y}(t) - \hat{\mathbf{y}}(t|t-1)]. \end{aligned}$$

In other words, the prediction error of $\mathbf{z}(t)$ is identical to the prediction error of $\mathbf{y}(t)$, except for a constant factor equal to the instantaneous gain of the filter. Therefore, not only is fitting a model by minimizing the prediction error of $\mathbf{z}(t)$ identical to fitting a model by minimizing the prediction error of $\mathbf{y}(t)$, but the cross-validated R^2 (up to a fixed constant) and residual whiteness of these models are also identical.

This argument clearly fails for nonlinear systems. Indeed, fitting a nonlinear model on $\mathbf{z}(t)$ can result in a different model with different R^2 and residual whiteness than the same model fit on $\mathbf{y}(t)$. The critical point, however, is that the linear filter $F(q)$ cannot generate nonlinearity, while it can certainly weaken and even eliminate it. We explain these two points in more detail next.

2. Assume, first, that the dynamics of $\mathbf{y}(t)$ are truly linear, as seems to be the case from our analysis in the main text. Then, the relationship between $\Delta\mathbf{y}(t)$ and $\mathbf{y}(t-1)$ (as random vectors) is linear. By definition, if the relationship between two random vectors \mathbf{u} and \mathbf{v} is linear, they can be written in the form

$$\begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{e}_1 \\ \mathbf{e}_2 \end{bmatrix} \quad (\text{S1})$$

where \mathbf{A} is an appropriate matrix and \mathbf{e}_1 and \mathbf{e}_2 are independent. Now let \mathbf{u}_F be the result of applying a linear filter $F(q)$ to (the samples of) \mathbf{u} . In other words, if N is the order of an (arbitrarily accurate) FIR approximation of $F(q)$,

$$\mathbf{u}_F = [\mathbf{u}_1 \quad \mathbf{u}_2 \quad \cdots \quad \mathbf{u}_N] \begin{bmatrix} f(0) \\ f(1) \\ \vdots \\ f(N-1) \end{bmatrix} = [\mathbf{u}_1 \quad \mathbf{u}_2 \quad \cdots \quad \mathbf{u}_N] \mathbf{f} \quad (\text{S2})$$

where $f(t)$ is the impulse response of $F(q)$ and $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N$ are identically distributed (but not necessarily independent) samples of \mathbf{u} . From Eq. (S1),

$$[\mathbf{u}_1 \quad \mathbf{u}_2 \quad \cdots \quad \mathbf{u}_N] = [\mathbf{A}_{11} \quad \mathbf{A}_{12}] \begin{bmatrix} \mathbf{e}_{1,1} & \mathbf{e}_{1,2} & \cdots & \mathbf{e}_{1,N} \\ \mathbf{e}_{2,1} & \mathbf{e}_{2,2} & \cdots & \mathbf{e}_{2,N} \end{bmatrix} \quad (\text{S3})$$

where $\mathbf{e}_{1,1}, \dots, \mathbf{e}_{1,N}$ are identically distributed (but not necessarily independent) samples of \mathbf{e}_1 , similarly for \mathbf{e}_2 . Note that each $\mathbf{e}_{1,t}$ is still independent from each $\mathbf{e}_{2,s}$ by definition. Substituting Eq. (S3) into Eq. (S2) thus gives

$$\begin{aligned} \mathbf{u}_F &= [\mathbf{A}_{11} \quad \mathbf{A}_{12}] \begin{bmatrix} \mathbf{e}_{1,1} & \mathbf{e}_{1,2} & \cdots & \mathbf{e}_{1,N} \\ \mathbf{e}_{2,1} & \mathbf{e}_{2,2} & \cdots & \mathbf{e}_{2,N} \end{bmatrix} \mathbf{f} \\ &= [\mathbf{A}_{11} \quad \mathbf{A}_{12}] \begin{bmatrix} \mathbf{e}_{1,F} \\ \mathbf{e}_{2,F} \end{bmatrix} \end{aligned}$$

where $\mathbf{e}_{1,F}$ and $\mathbf{e}_{2,F}$ are filtered versions of \mathbf{e}_1 and \mathbf{e}_2 and, still, independent of each other. Following the same steps for \mathbf{v} , we get

$$\begin{bmatrix} \mathbf{u}_F \\ \mathbf{v}_F \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{e}_{1,F} \\ \mathbf{e}_{2,F} \end{bmatrix}$$

showing that \mathbf{u}_F and \mathbf{v}_F are still linearly related after filtering. This is indeed expected since a linear filter cannot generate nonlinear dependence between two signals that are originally linearly related.

3. Now assume, in contrast, that the dynamics of $\mathbf{y}(t)$ is in fact nonlinear and we filter $\mathbf{y}(t)$ to get $\mathbf{z}(t) = F(q)\mathbf{y}(t)$. The best that can happen, as far as detecting nonlinearity is concerned, is that the dynamics of $\mathbf{z}(t)$ remain nonlinear. However, it is possible that $F(q)$ weakens or completely averages out the nonlinearities in $\mathbf{y}(t)$, as we saw in the main text. In fact, the common bandpass filter over $[0.01, 0.08]$ Hz (compared to a Nyquist frequency of $1/2\text{TR} \simeq 0.7$ Hz) is strongly lowpass and involves significant averaging over time.

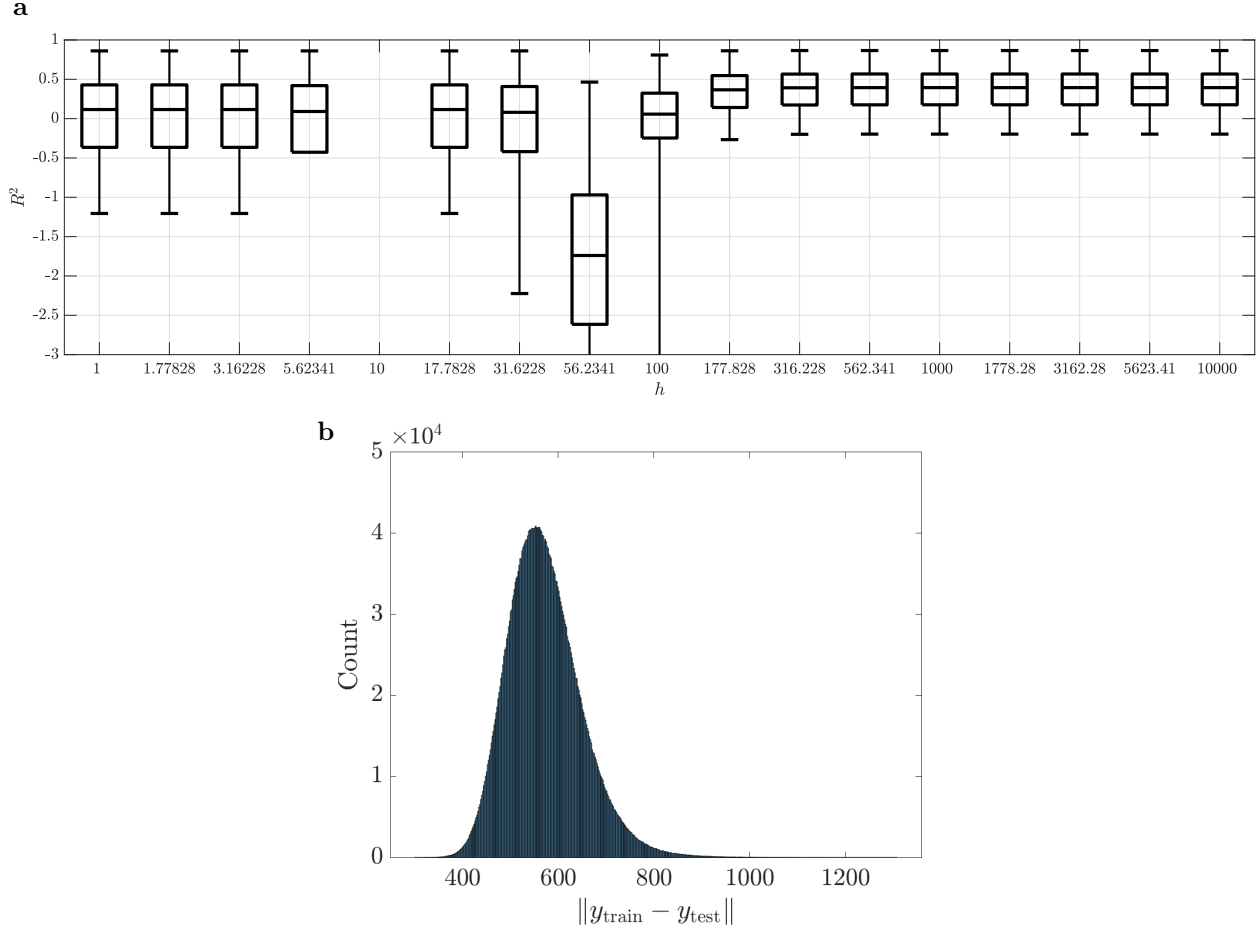
In conclusion, while bandpass filtering has no effect on linear models and preserves linearity of time series, it can well weaken/eliminate any nonlinearity in the time series. Therefore, regardless of how much “cleaner” bandpass-filtered data might be, finding no nonlinearity before bandpass filtering, as pursued in the main text, is a stronger statement than the same finding would be if obtained after bandpass filtering.

2 Supplementary Table 1: List of linear and nonlinear families of models

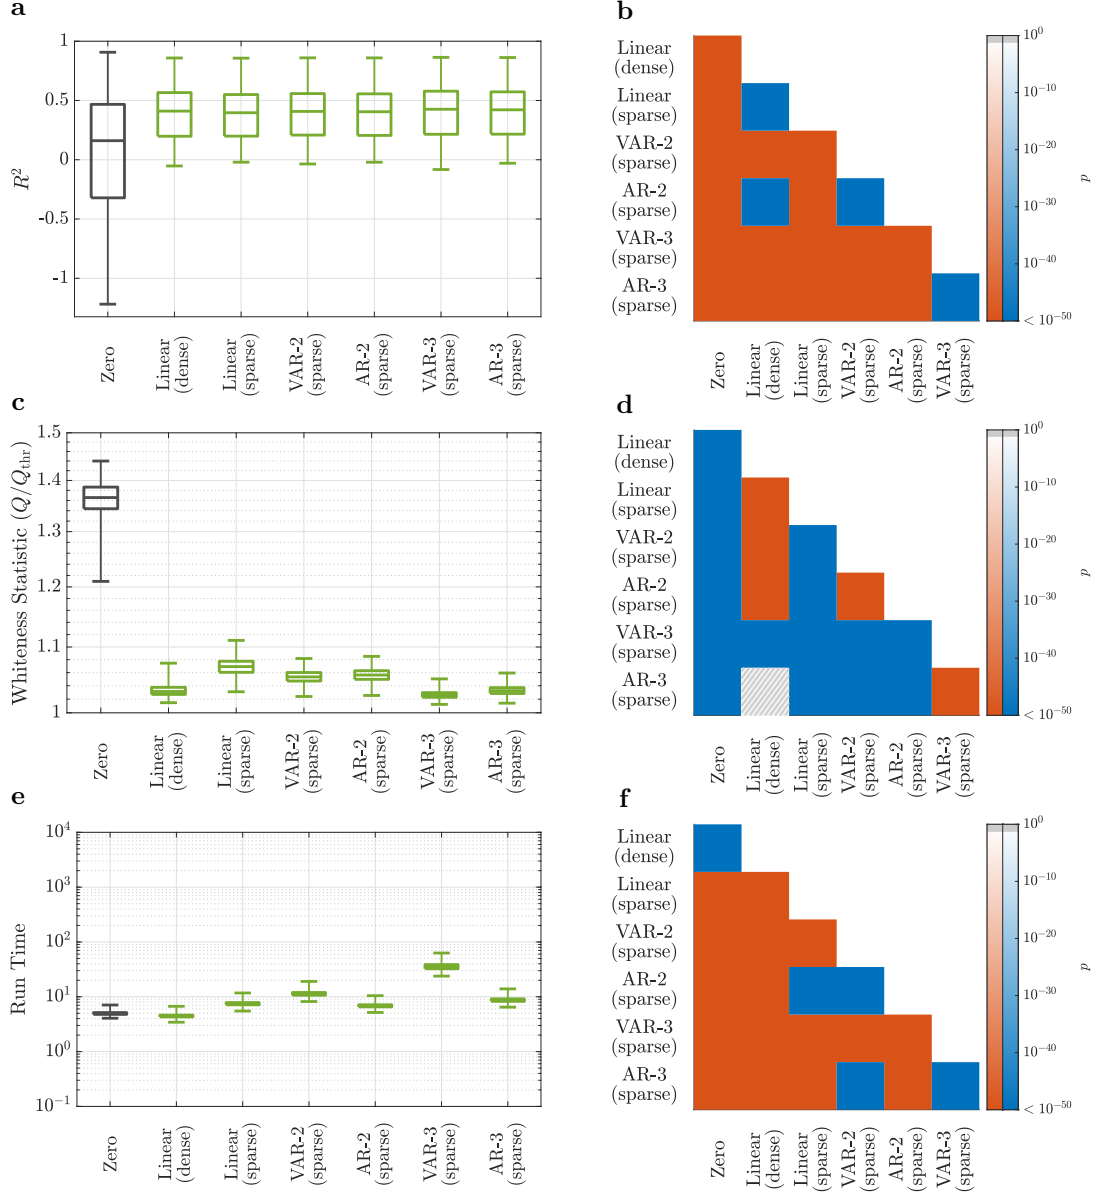
Supplementary Table 1 | List of linear and nonlinear families of models. The marks [†] and [‡] indicate, respectively, that a method is used only for fMRI or iEEG. See Methods for a description of each model.

Label	Title	Equation	Hyper-parameters
Linear (dense)	Linear models with states at the BOLD/LFP level	$\mathbf{y}(t) - \mathbf{y}(t-1) = \mathbf{W}\mathbf{y}(t-1) + \mathbf{e}(t)$	None
Linear (sparse)			$\lambda = 0.95$ (fMRI) $\lambda = 1.35$ (iEEG)
Linear (pairwise) [†]		$y_i(t) - y_i(t-1) = w_{ij}y_j(t-1) + e_i(t), i, j = 1, \dots, n$	None
AR-2 (sparse) [†]	Linear autoregressive models	$\mathbf{y}(t) - \mathbf{y}(t-1) = \mathbf{W}\mathbf{y}(t-1) + \mathbf{D}_2\mathbf{y}(t-2) + \mathbf{D}_3\mathbf{y}(t-3) + \dots + \mathbf{D}_d\mathbf{y}(t-d) + \mathbf{e}(t)$	$d = 2, \lambda = 0.95$, diagonal \mathbf{D}_2
VAR-2 (sparse) [†]			$d = 2, \lambda = 0.9$
AR-3 (sparse) [†]			$d = 3, \lambda = 0.5$, diagonal $\mathbf{D}_2, \mathbf{D}_3$
VAR-3 (sparse) [†]			$d = 3, \lambda = 0.35$
AR-100 (sparse) [‡]			$d = 112, \lambda = 1.35$
AR-100 (scalar) [‡]			$d = 112$
Linear w/ HRF [†]	Linear models with states at the neural level	$\mathbf{x}(t) - \mathbf{x}(t-1) = \mathbf{W}\mathbf{x}(t-1) + \mathcal{G}_1(q)\hat{\mathbf{e}}_1(t)$	$n_h = n_\phi = n_\psi = 5, \lambda = 11$
		$\mathbf{y}(t) = \mathcal{H}(q)\mathbf{x}(t) + \mathcal{G}_2(q)\hat{\mathbf{e}}_2(t)$	
		$\mathcal{H}(q) = \sum_{p=1}^{n_h} \text{diag}(\mathbf{H}_{:,p})q^{-p}$	
		$\mathcal{F}_1(q) = \mathbf{I} - \mathcal{G}_1^{-1}(q) = \sum_{p=1}^{n_\phi} \text{diag}(\mathbf{\Phi}_{:,p})q^{-p}$	
Subspace	Linear models with abstract data-driven states	$\mathcal{F}_2(q) = \mathbf{I} - \mathcal{G}_2^{-1}(q) = \sum_{p=1}^{n_\psi} \text{diag}(\mathbf{\Psi}_{:,p})q^{-p}$	$s = 1, r = 3, n_x = 25$ (fMRI) $s = 10, r = 69, n_x = 445$ (iEEG)
		$\mathbf{x}(t) - \mathbf{x}(t-1) = \mathbf{W}\mathbf{x}(t-1) + \mathbf{e}_1(t)$	
		$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t) + \mathbf{e}_2(t)$	
NMM	Nonlinear neural mass models	$\text{Cov} \left(\begin{bmatrix} \mathbf{e}_1(t) \\ \mathbf{e}_2(t) \end{bmatrix} \right) = \begin{bmatrix} \mathbf{Q} & \mathbf{M} \\ \mathbf{M}^T & \mathbf{R} \end{bmatrix}$	MINDy default (fMRI) $\lambda_1 = \lambda_2 = 0.2, \lambda_3 = 0.7, \lambda_4 = 0.5$ (iEEG)
		$\mathbf{y}(t) - \mathbf{y}(t-1) = (\mathbf{W}\psi_\alpha(\mathbf{y}(t-1)) - \mathbf{D}\mathbf{y}(t-1))\Delta_T + \mathbf{e}(t)$	
NMM w/ HRF [†]		$\mathbf{x}(t) - \mathbf{x}(t-1) = (\mathbf{W}\psi_\alpha(\mathbf{x}(t-1)) - \mathbf{D}\mathbf{x}(t-1))\Delta_T + \mathbf{e}_1(t)$	MINDy default
		$\mathbf{y}(t) = \mathcal{H}(q)\mathbf{x}(t) + \mathbf{e}_2(t)$	
DNN (MLP)	Nonlinear models via multi-layer perceptron deep neural networks	$\mathbf{y}(t) - \mathbf{y}(t-1) = f(\mathbf{y}(t-1), \dots, \mathbf{y}(t-d)) + \mathbf{e}(t)$	$d = 1, D = 6, W = 2$ (fMRI) $d = 4, D = 3, W = 29$ (iEEG)
DNN (CNN)	Nonlinear models via convolutional deep neural networks	$\mathbf{y}(t) - \mathbf{y}(t-1) = f(\mathbf{y}(t-1), \dots, \mathbf{y}(t-d)) + \mathbf{e}(t)$	$d = 17, D = 2, l_{\text{filt}} = 7, n_{\text{filt}} = 11, n_{\text{pool}} = 4, p_{\text{drop}} = 0.4$ (fMRI) $d = 13, D = 2, l_{\text{filt}} = 5, n_{\text{filt}} = 10, n_{\text{pool}} = 2, p_{\text{drop}} = 0.51$ (iEEG)
LSTM (IIR)	Nonlinear models via long short-term	$\mathbf{y}(t) - \mathbf{y}(t-1) = f(\mathbf{y}(t-1), \dots, \mathbf{y}(0)) + \mathbf{e}(t)$	$W = 12$ (fMRI) $W = 7$ (iEEG)
LSTM (FIR)	memory recurrent neural networks	$\mathbf{y}(t) - \mathbf{y}(t-1) = f(\mathbf{y}(t-1), \dots, \mathbf{y}(t-d)) + \mathbf{e}(t)$	$d = 1, W = 16$ (fMRI) $d = 32, W = 2$ (iEEG)
Manifold	Nonlinear manifold-based models	$\mathbf{y}(t) - \mathbf{y}(t-1) = f(\mathbf{y}(t-1), \dots, \mathbf{y}(t-d)) + \mathbf{e}(t)$	$d = 1, h = 830$ (fMRI) $d = 7, h = 1.3 \times 10^4$ (iEEG)
MMSE (pairwise) [†]	Nonlinear minimum mean squared error models (optimal)	$y_i(t) - y_i(t-1) = \mathbb{E}[y_i(t) - y_i(t-1) y_j(t-1)], i, j = 1, \dots, n$	$N = 280, \beta = 0.156$
MMSE (scalar) [‡]		$y_i(t) - y_i(t-1) = \mathbb{E}[y_i(t) - y_i(t-1) y_i(t-1), \dots, y_i(t-d)], i = 1, \dots, n$	$d = 9, N = 309, \beta = 0.007$
Zero	Zero model	$\mathbf{y}(t) - \mathbf{y}(t-1) = \mathbf{e}(t)$	None

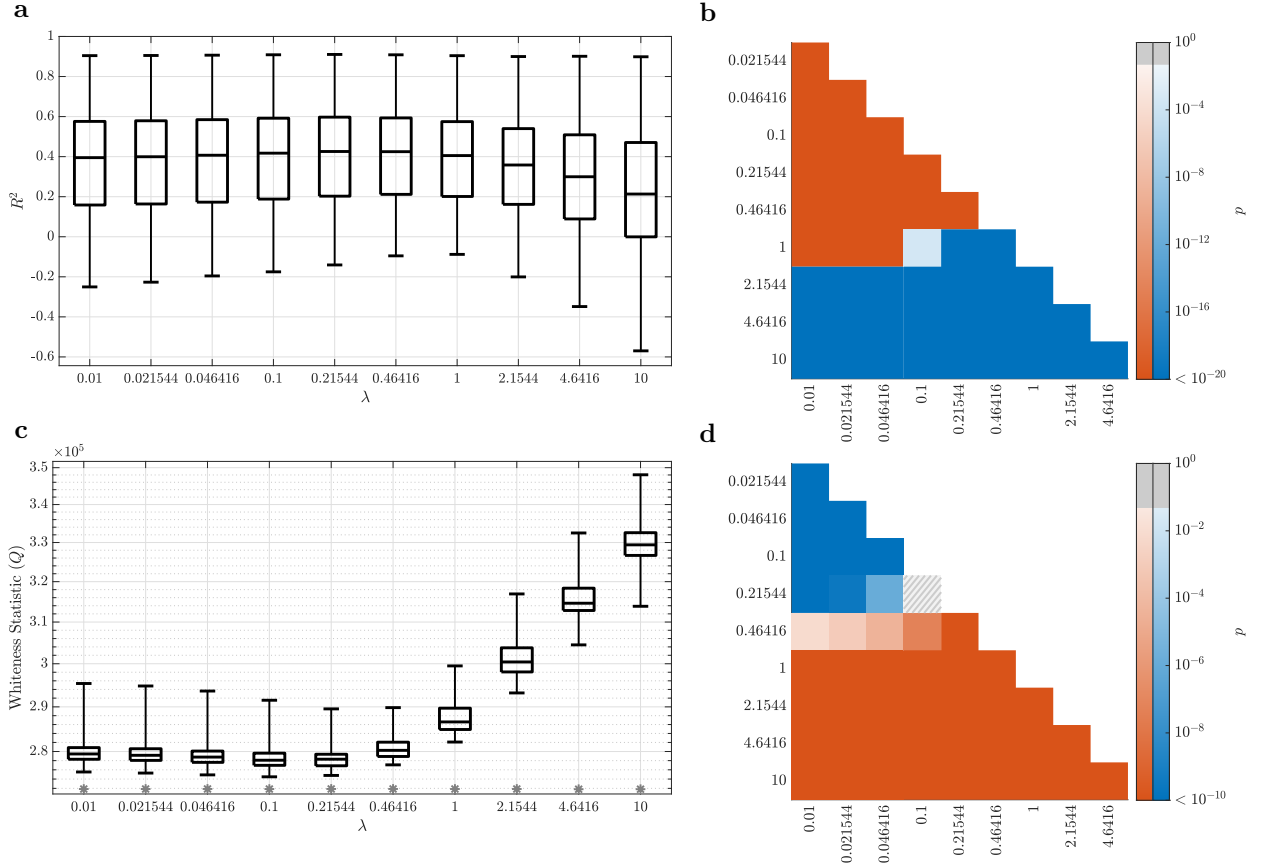
3 Supplementary Figures



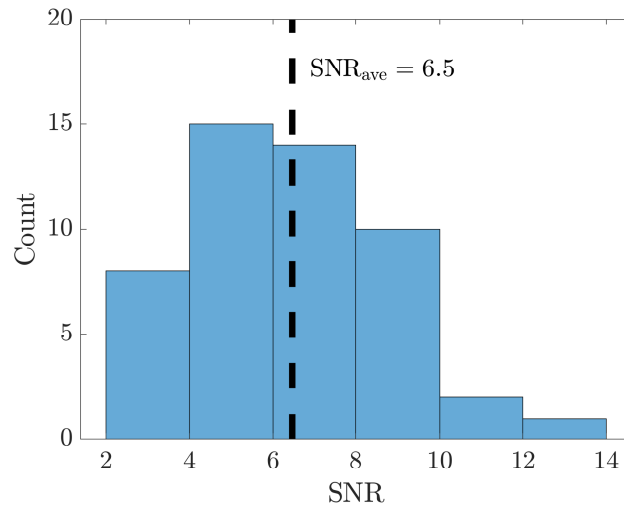
Supplementary Fig. 1 | Effect of window size h on the accuracy of the manifold-based locally linear ('Manifold') model in fMRI data. (a) Boxplot of the R^2 distribution as a function of h , combined across the 116 brain regions of 70 randomly selected subjects (10% of all subjects to reduce computational cost). The R^2 values were not computable ('NaN') for $h = 10$ due to limited machine precision. The model is equivalent to the zero model for the three leftmost boxplots as no training point falls within the Gaussian-weighted neighborhood of any test points. As h is increased to $h \sim 10$, few training data points start to fall within the neighborhood window of some of the test points, but are far enough that their Gaussian weights fall below machine precision, leading to missing ('NaN') predicted values and, hence, R^2 . As h is further increased, more training points fall within the neighborhood of each test point, but are few enough to lead to poor R^2 , until h is increased enough to reach the globally linear regime. (c) The distribution of the Euclidean distance between all pairs of training and test points for a randomly selected subject (subject 103818) to aid in understanding the trends that are apparent in panels (a) and (b). In particular, note that $h = 10^4$ (and even smaller values) clearly lead to a globally linear model as almost all of the training-test pairs of points have distances less than $h/10$.



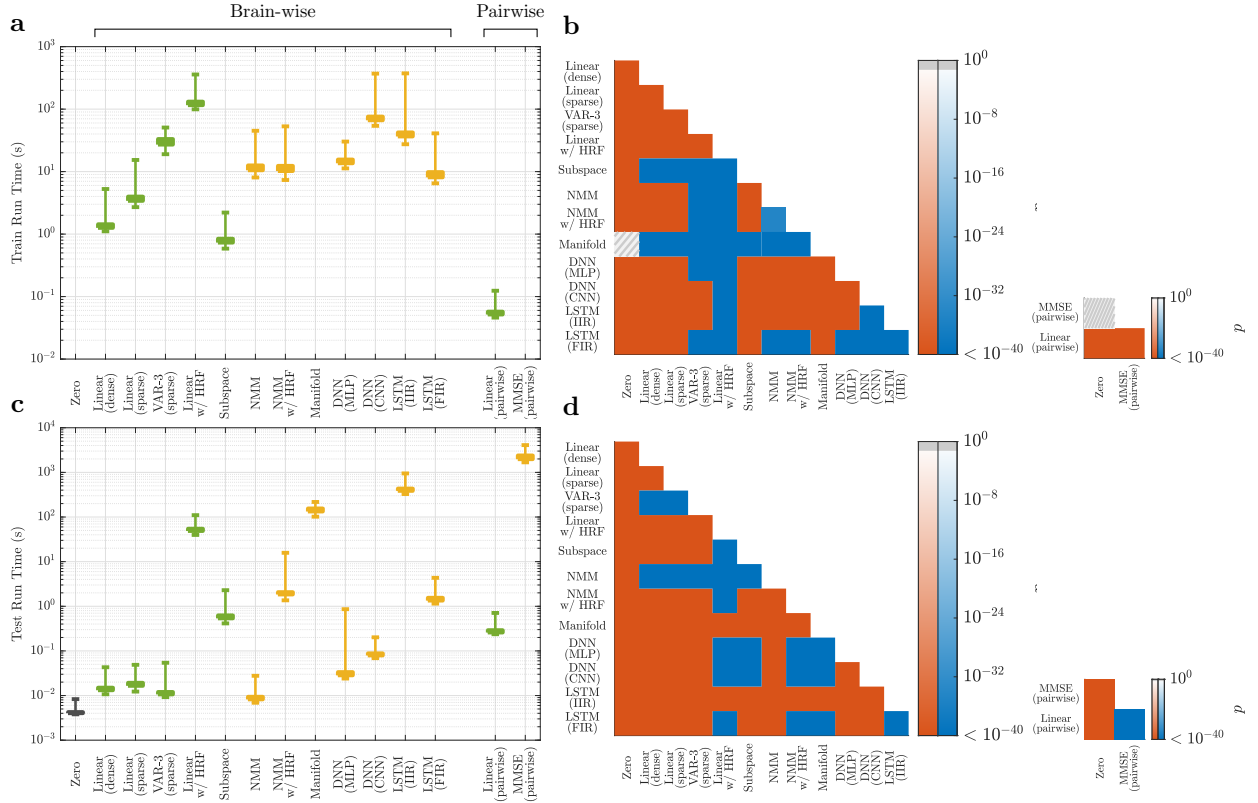
Supplementary Fig. 2 | The effects of the number of lags and sparsity patterns on the prediction accuracy and computational complexity of linear AR models of rsfMRI. Panels parallel those in **Fig. 2** in the main text and the descriptions of method acronyms are given in **Table 1** therein. Generally, the number of lags and sparsity patterns have little effect on the prediction accuracy of linear AR models for rsfMRI data (in contrast to rsiEEG data, as explained in the main text), as seen from panel (a). The estimates of statistical significance in panel (b) are to a great extent due to the large sample size (700×116). However, increasing the number of regressors, both by increasing the number of AR lags and by allowing for off-diagonal entries of all lags (‘VAR’ models), does lead to a non-trivial improvement in the whiteness of residuals, even though is often accompanied by non-trivial increases in model complexity and computation time as well. In all box plots, the center line, box limits, and whiskers represent the median, upper and lower quartiles, and the smallest and largest samples, respectively, and sample size = 81200.



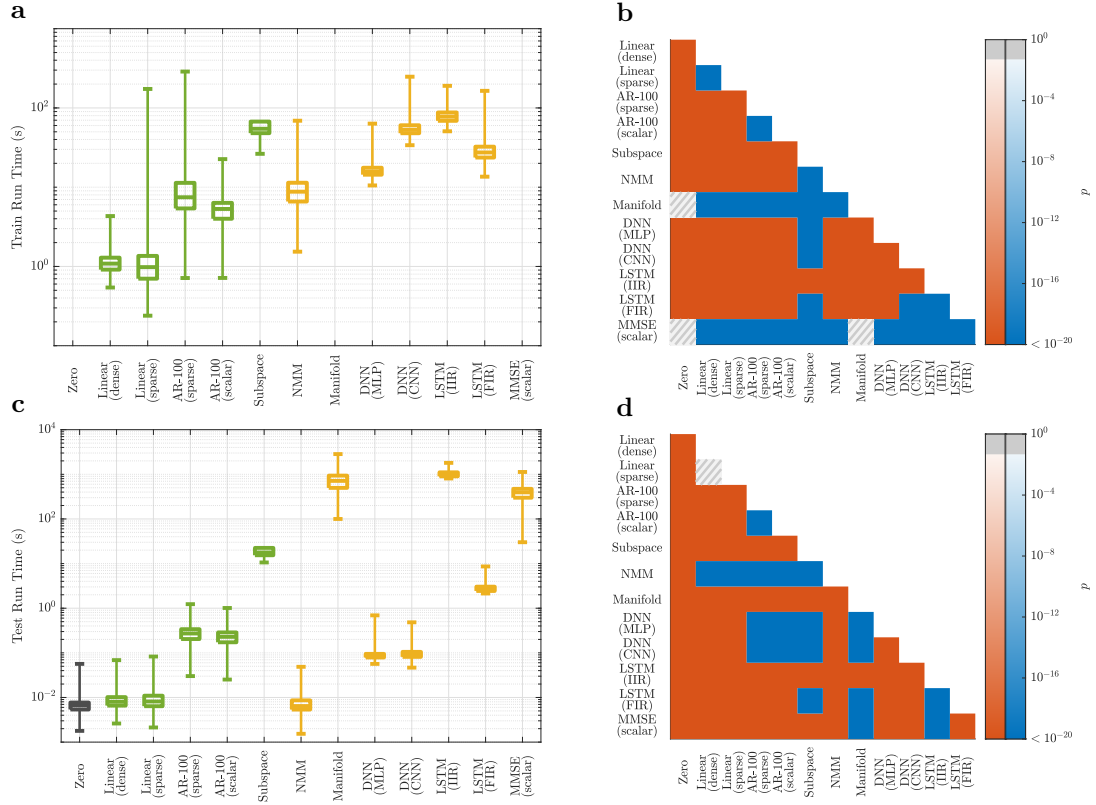
Supplementary Fig. 3 | The effect of the LASSO parameter λ on the accuracy of the 'VAR-3 (sparse)' model. Panels parallel those in Fig. 2 of the main text. (a) The distribution of the cross-validated regional R^2 , combined across all regions and 10% of subjects (randomly selected), for varying values of λ . (b) The p -values of the one-sided Wilcoxon signed rank test performed between all pairs of distributions of R^2 in panel (a). (c, d) Similar to panels (a, b) but for the Q statistic of the test of whiteness of the residuals.



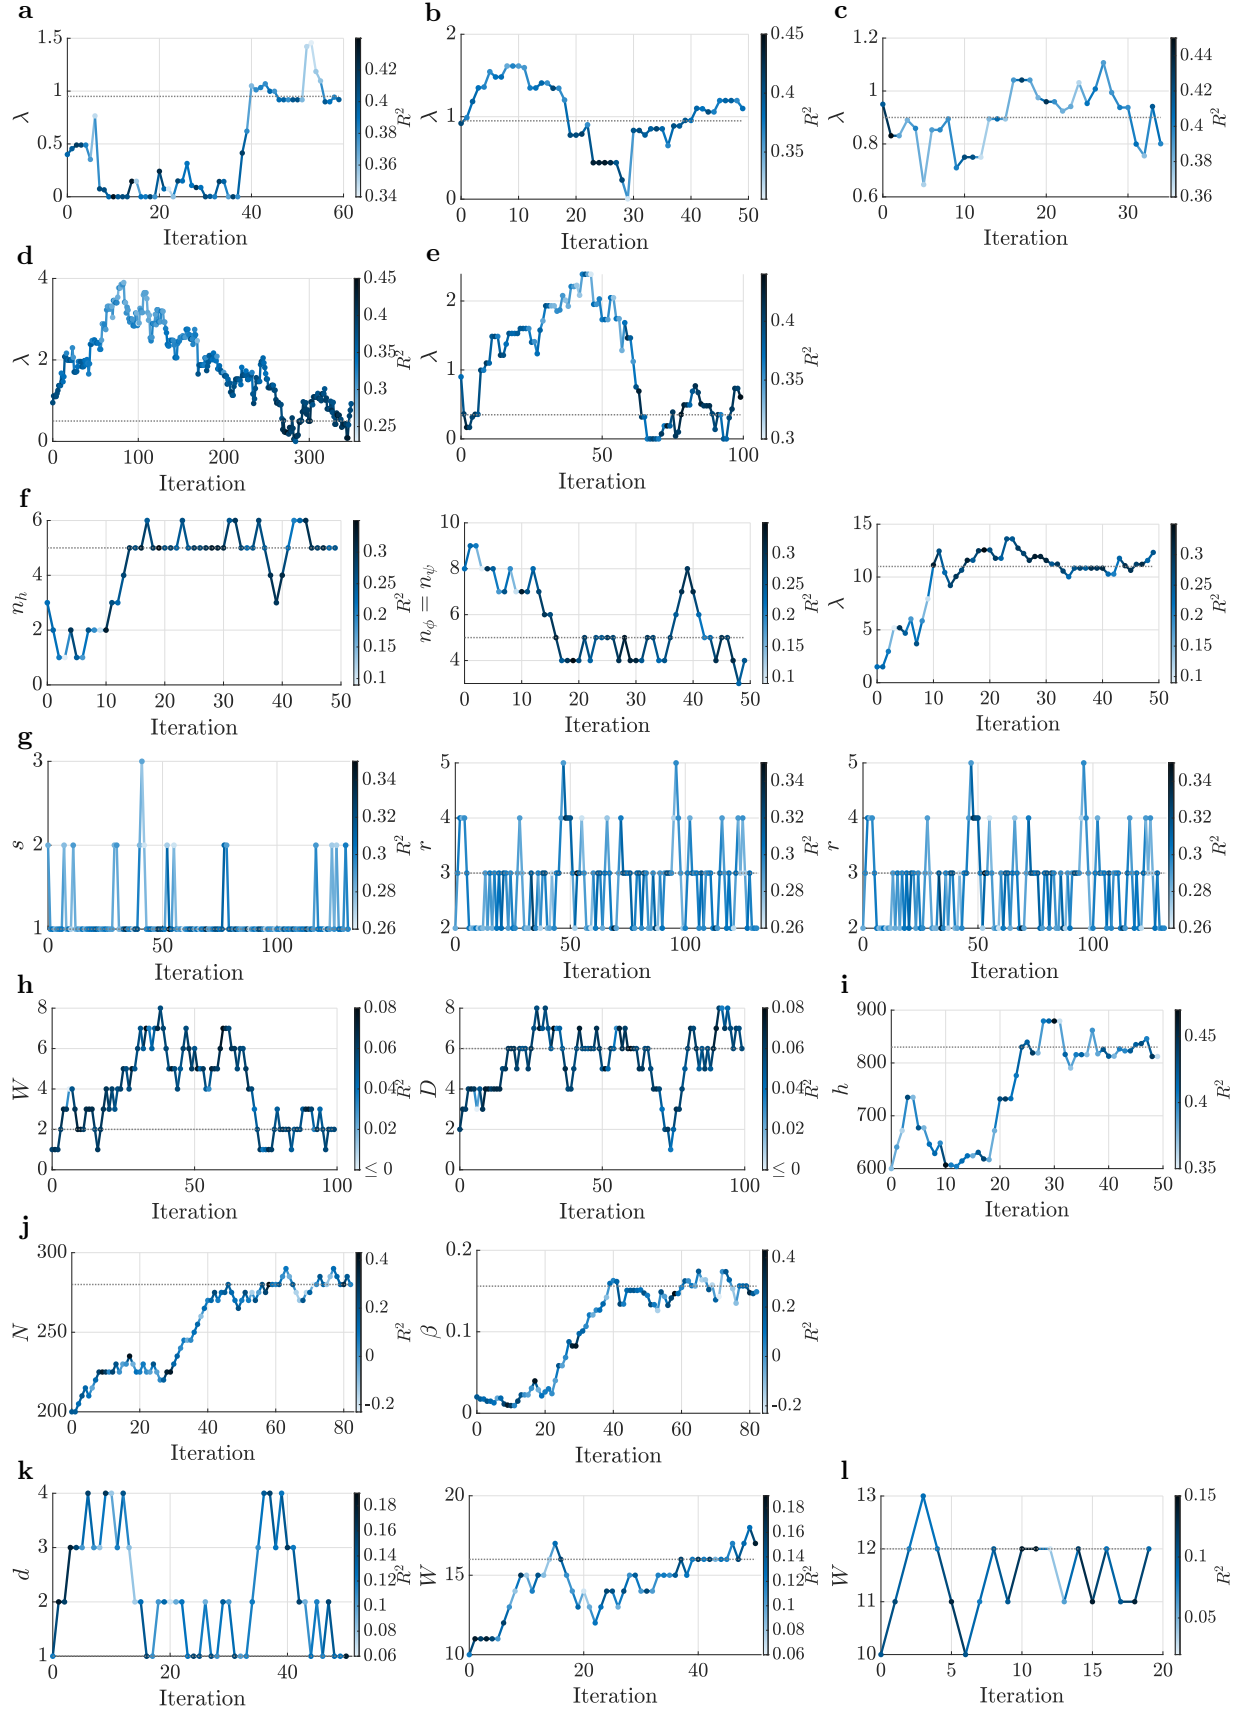
Supplementary Fig. 4 | Histogram of scanner SNR estimates for rsfMRI data. Scanner SNR was estimated for 50 randomly selected rest scans by comparing the average signal powers inside the respective subject’s gray matter and outside of their head (see Methods). Due to the conservatism of this method, the resulting SNR estimates are expectedly over-estimated, but yet are not far from the $\text{SNR} = 1$ level that is enough to completely mask nonlinear interactions on its own (**Fig. 4g-h**).

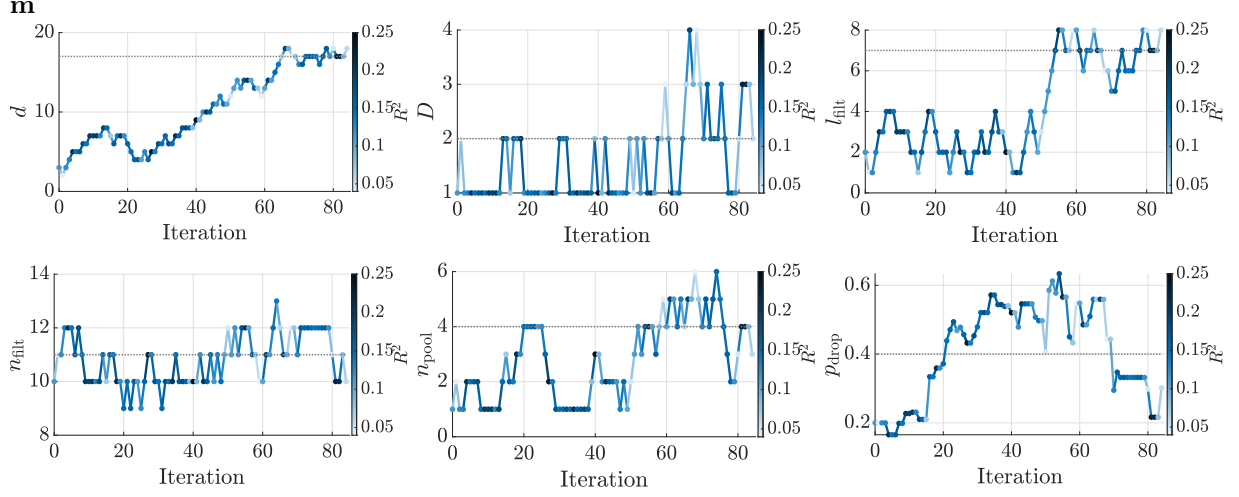


Supplementary Fig. 5 | Separate training and test times for system identification methods of rsfMRI. Details in panels (a,b) and (c,d) parallel those in of **Fig. 2e,f** in the main text. Note that the ‘Zero’, ‘Manifold’, and ‘MMSE (pairwise)’ methods do not have a training time by definition and hence we have set their training times uniformly to zero.



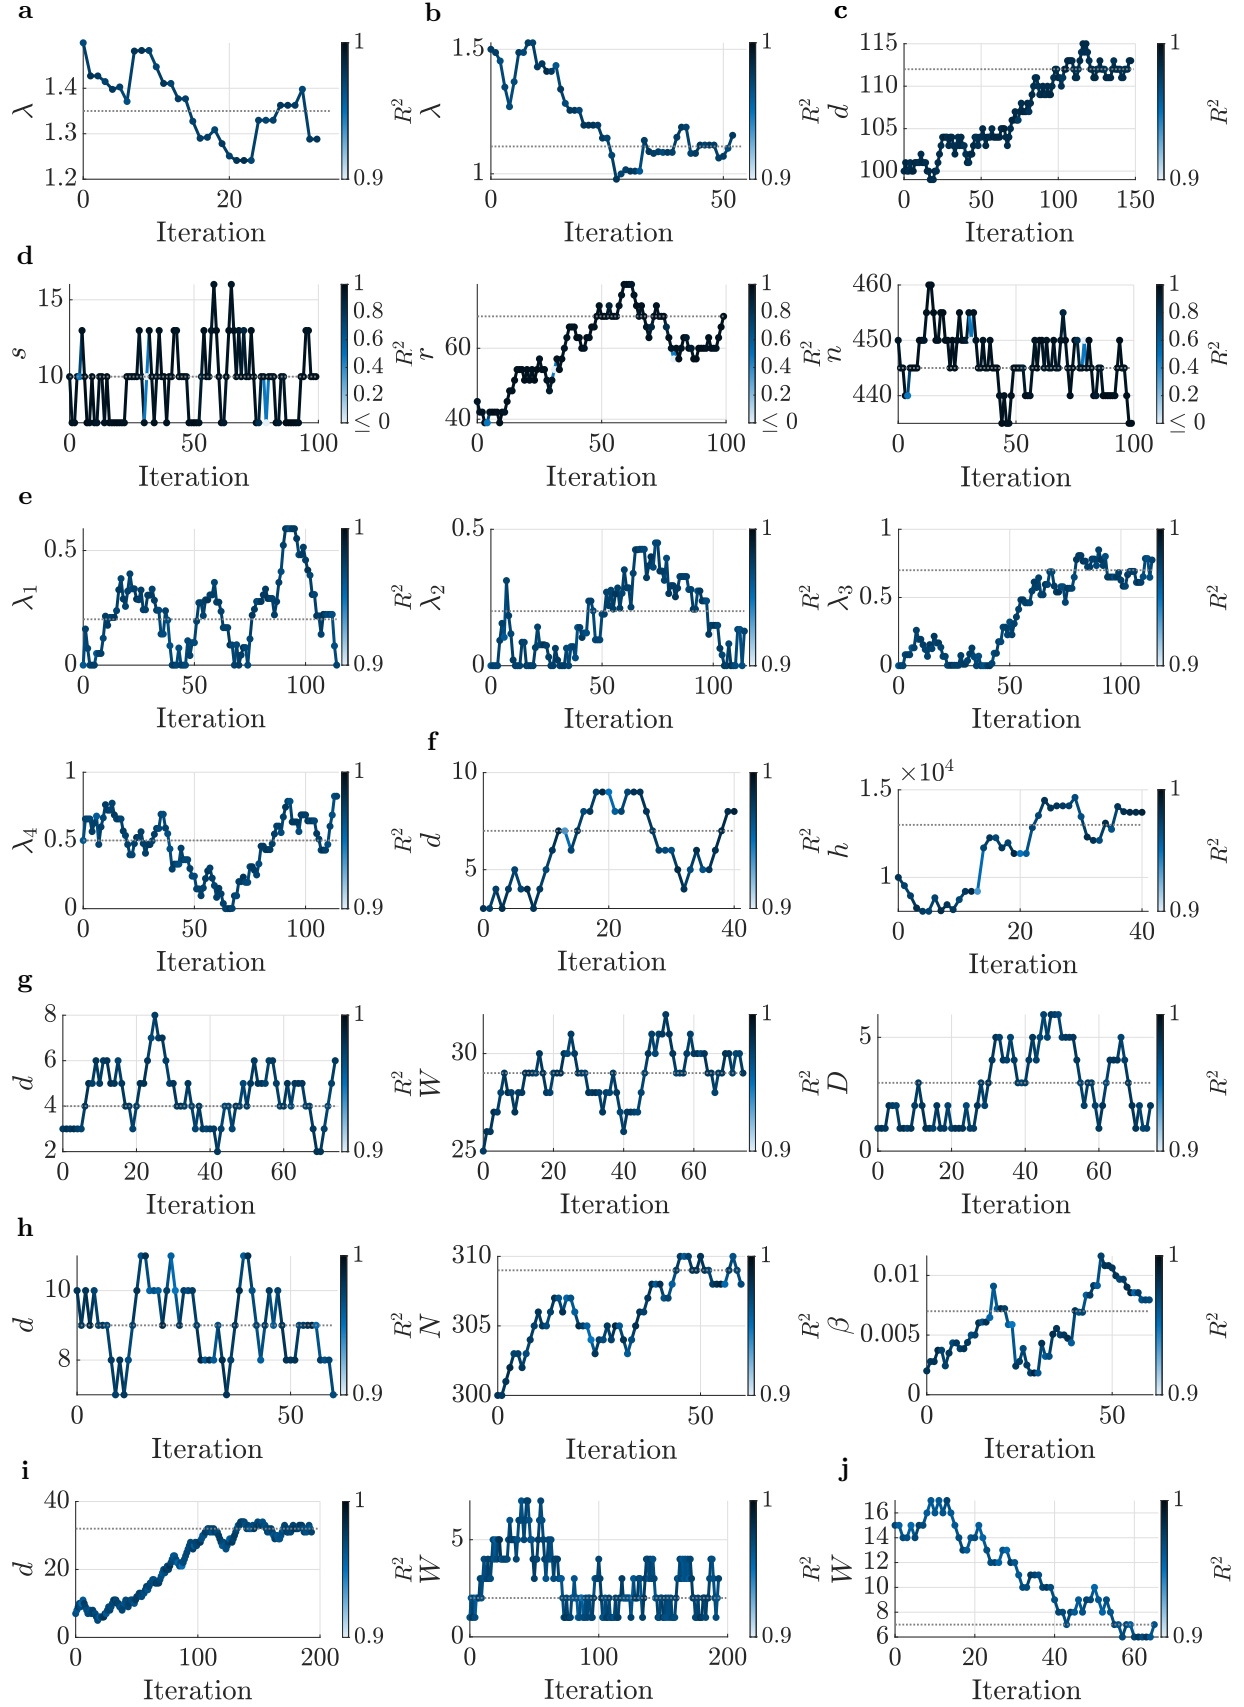
Supplementary Fig. 6 | Separate training and test times for system identification methods of rsiEEG. Details in panels (a,b) and (c,d) parallel those in of **Fig. 3e,f** in the main text. Note that the ‘Zero’, ‘Manifold’, and ‘MMSE (scalar)’ methods do not have a training time by definition and hence we have set their training times uniformly to zero.

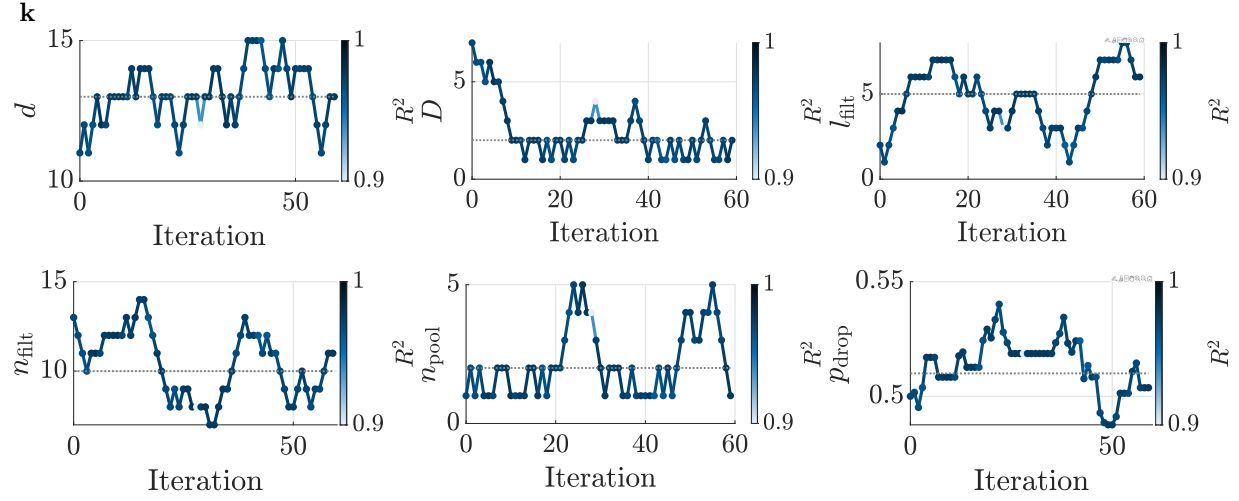




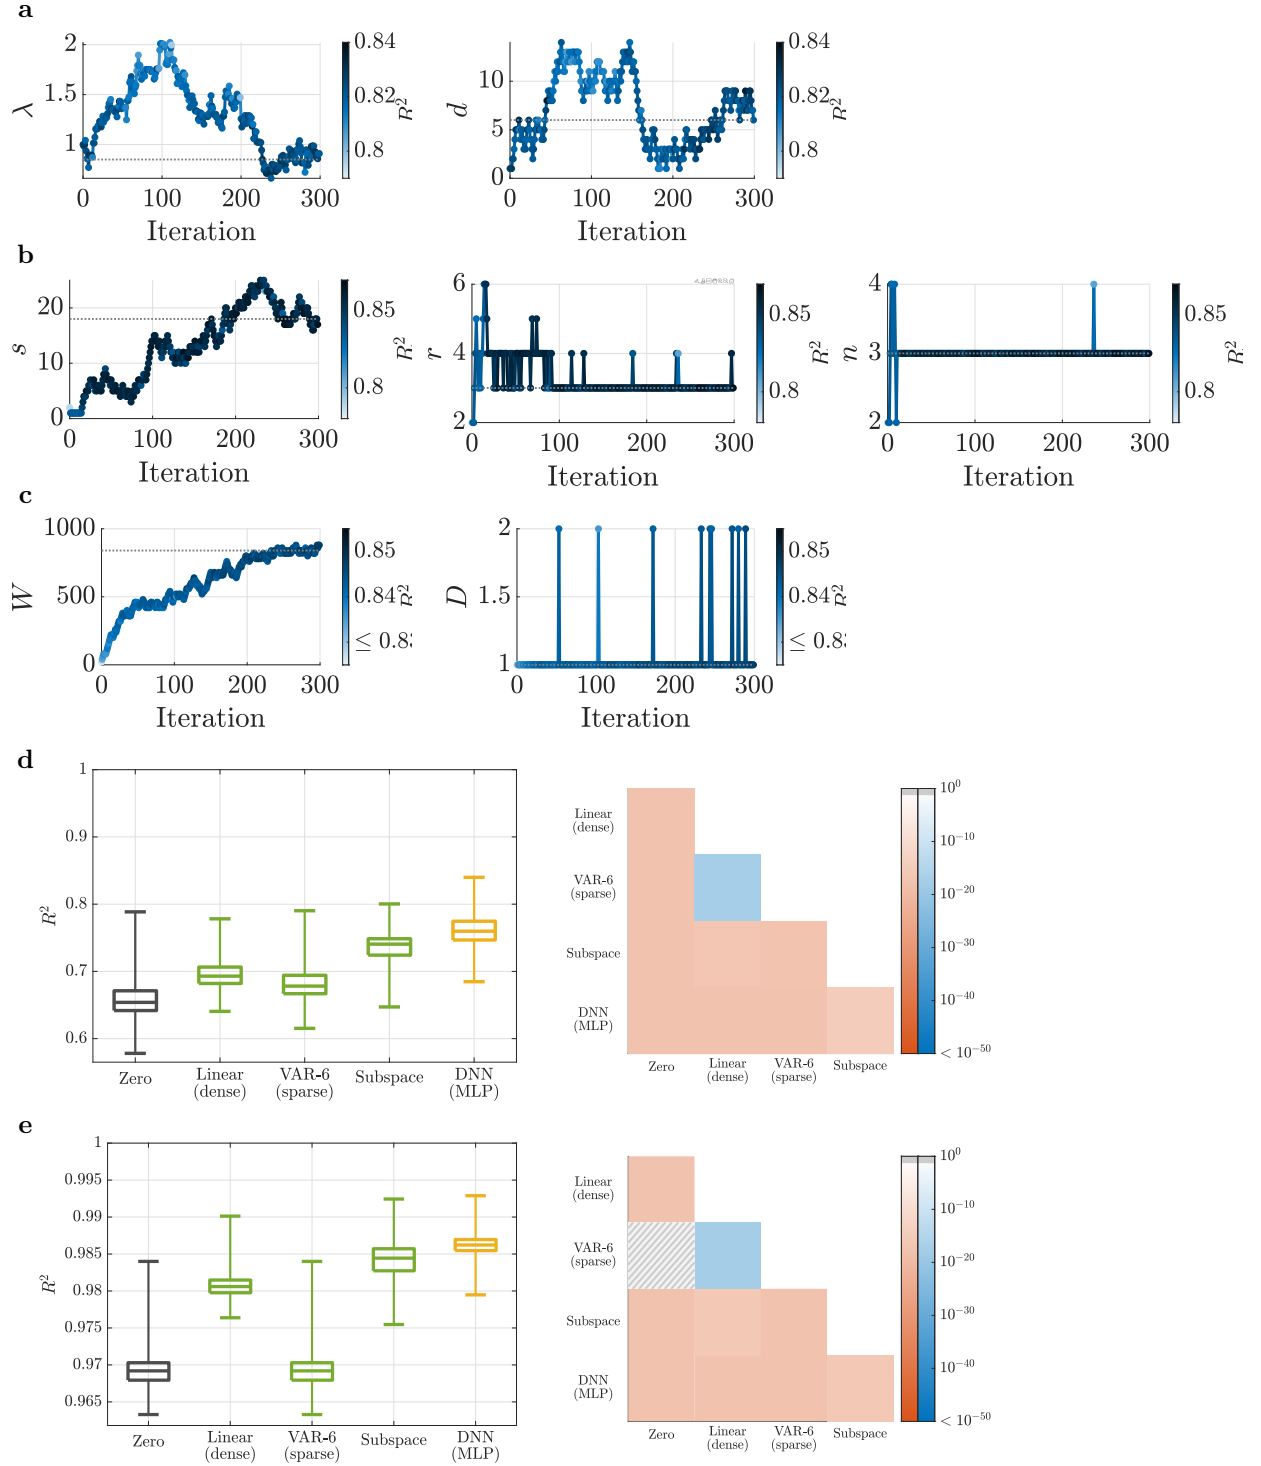
Supplementary Fig. 7 | Hyper-parameter tuning of linear and nonlinear model families for fMRI.

For each parametric family of models, its hyper-parameters were simultaneously optimized using stochastic gradient descent (SGD, see Methods) to select the model with the highest cross-validated R^2 within that model family. (a) Linear (sparse); (b) AR-2 (sparse); (c) VAR-2 (sparse); (d) AR-3 (sparse); (e) VAR-3 (sparse); (f) Linear w/ HRF; (g) Subspace; (h) DNN (MLP); (i) Manifold; (j) MMSE (pairwise); (k) LSTM (FIR); (l) LSTM (IIR); (m) DNN (CNN). Each panel shows the evolution of the hyper-parameter(s) of one model family during the SGD iterations, color-coded with the value of R^2 at each iteration, and the hyper-parameter value(s) selected as optimal (dotted gray lines, also given in **Table 1** in the main text). Note that the hyper-parameter(s) are not expected to converge to a fixed (optimal) value, but rather to fluctuate around it due to the stochastic nature of SGD and the natural variability of data segments.

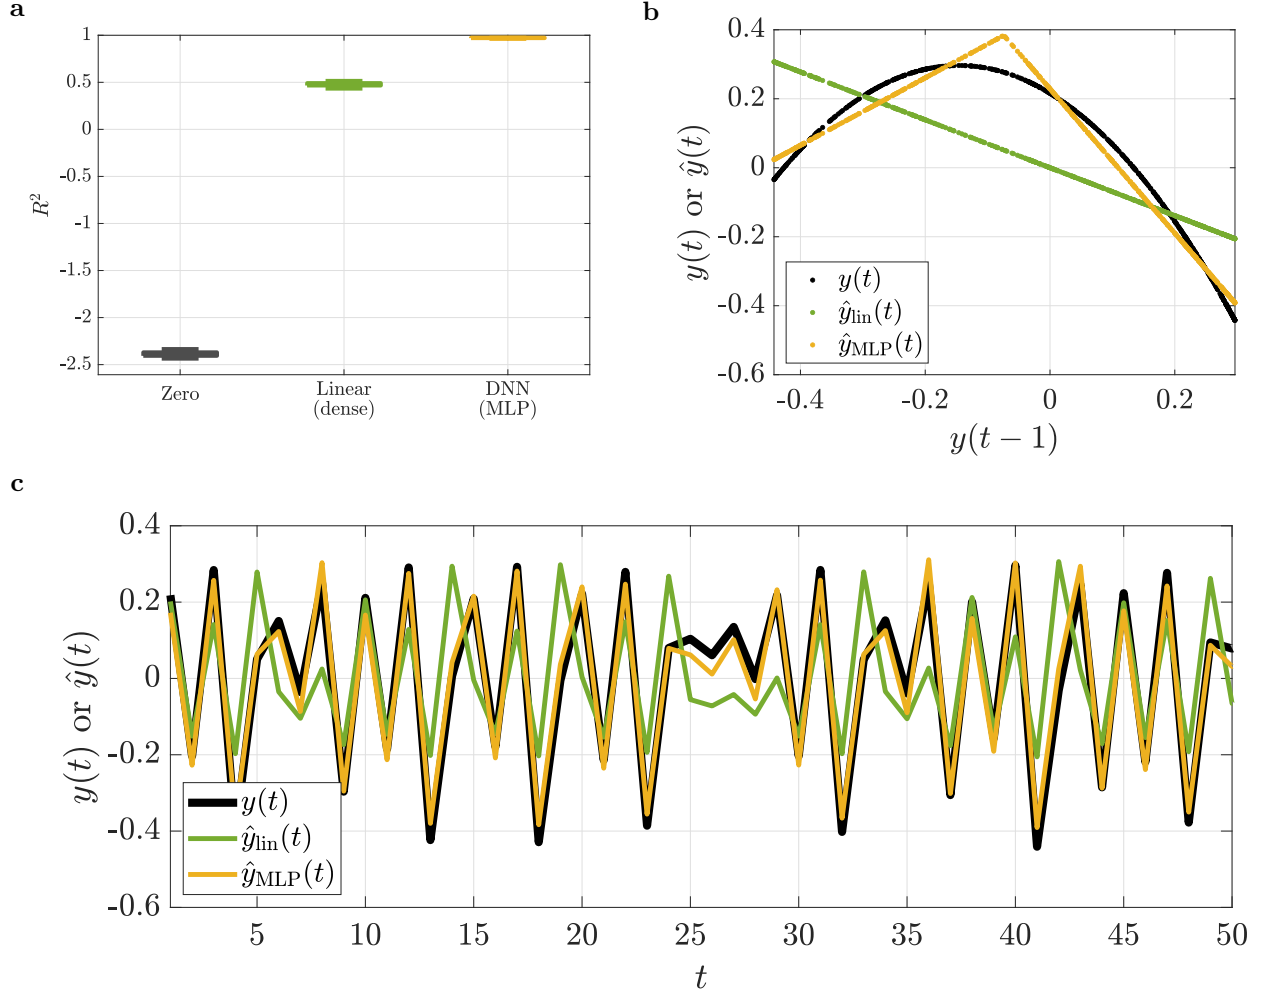




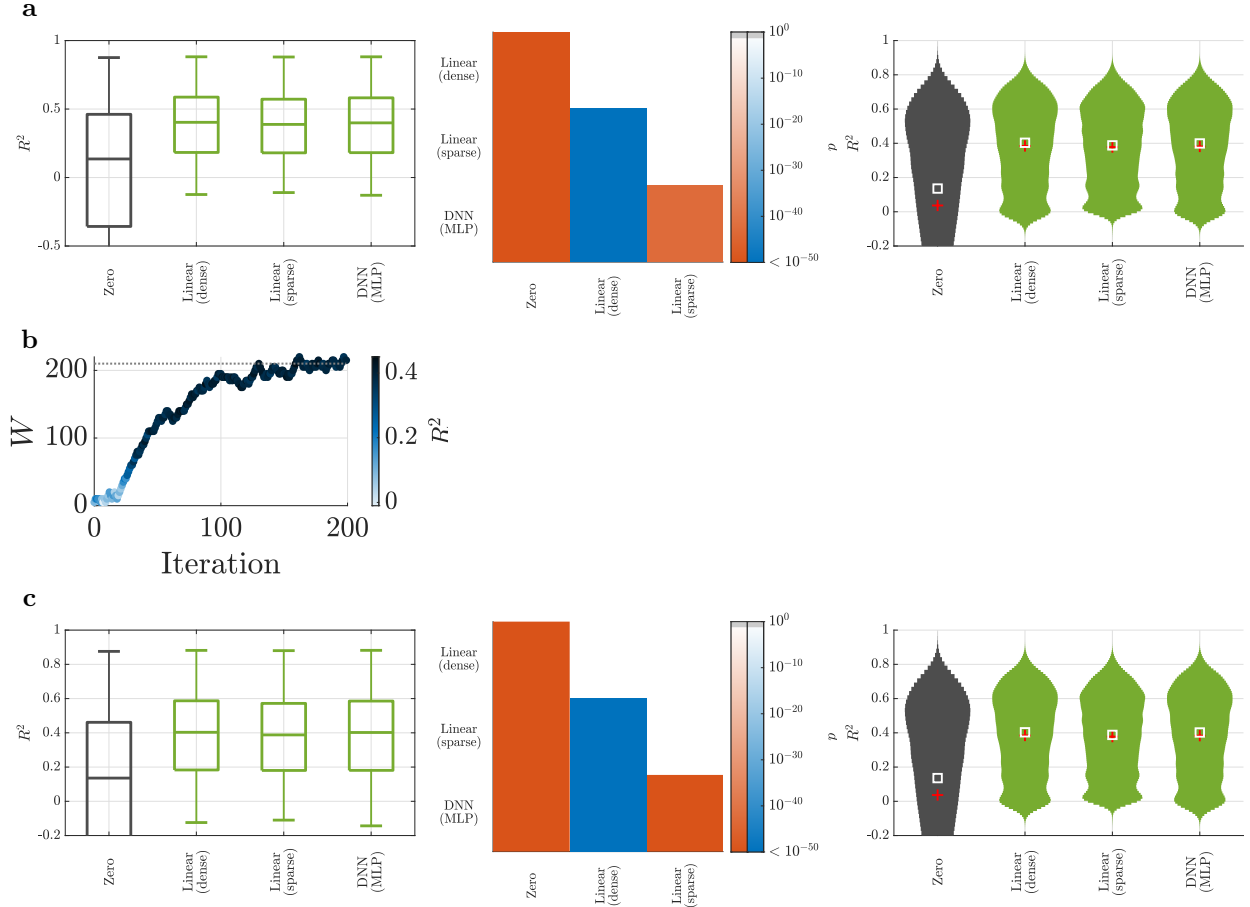
Supplementary Fig. 8 | Hyper-parameter tuning of linear and nonlinear model families for iEEG. (a) Linear (sparse); (b) AR-100 (sparse); (c) AR-100 (scalar); (d) Subspace; (e) NMM; (f) Manifold; (g) DNN; (h) MMSE (scalar); (i) LSTM (FIR); (j) LSTM (IIR); (k) DNN (CNN). Details and interpretations are the same as **Supplementary Fig. 7**.



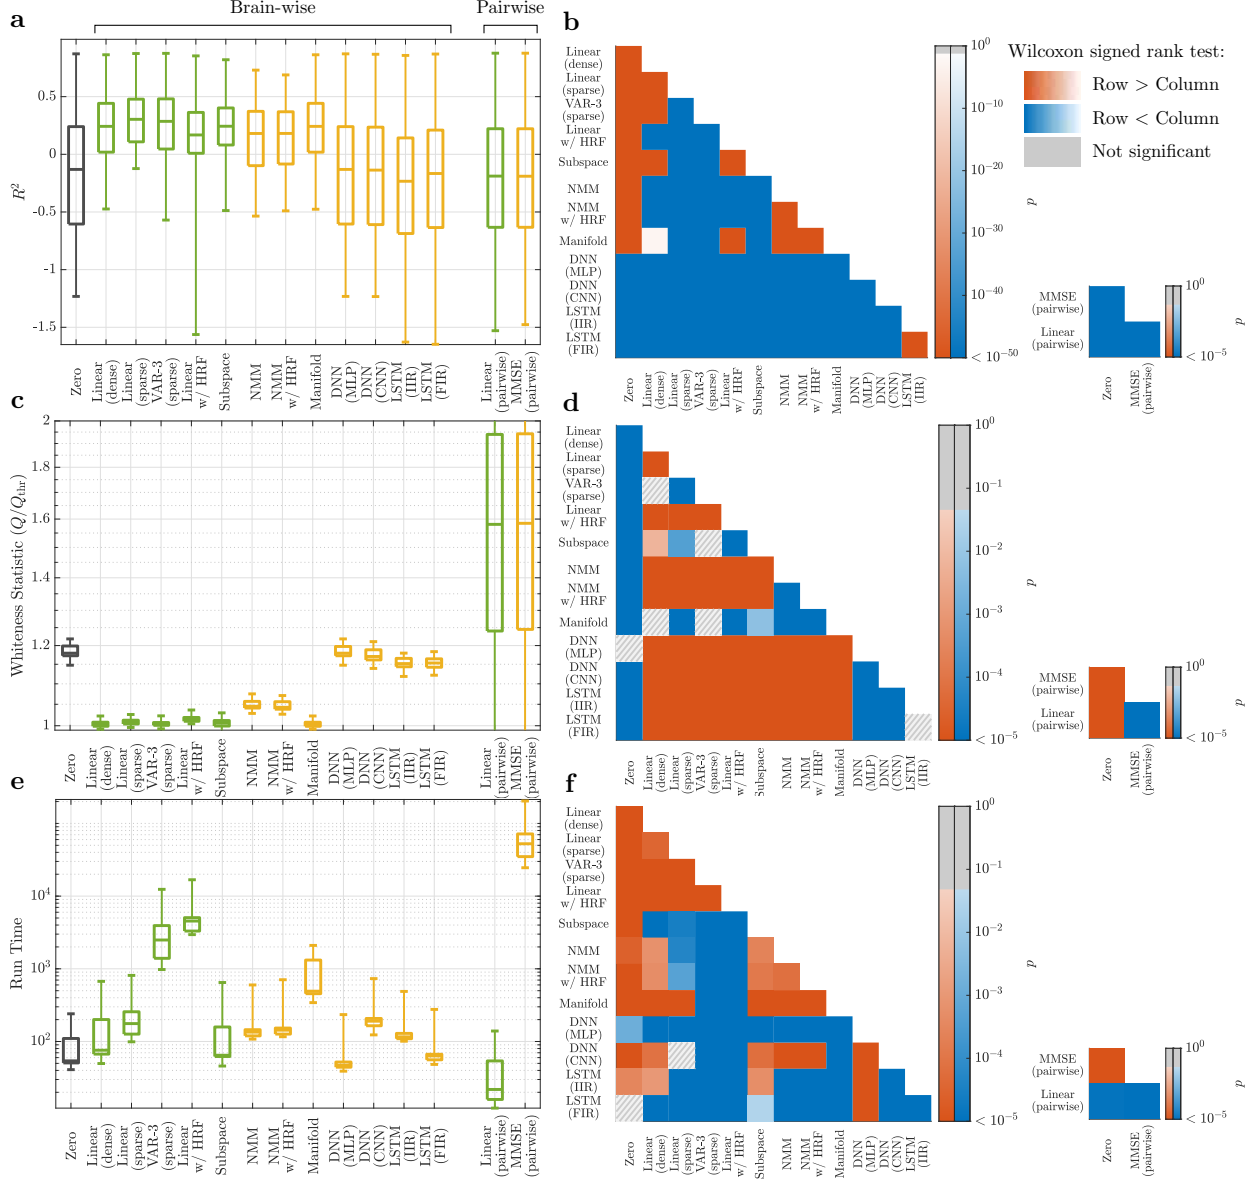
Supplementary Fig. 9 | Comparing the ‘DNN (MLP)’ model with linear models in fitting simulated data from the Izhikevic model. Panels (a-c) parallel those in **Supplementary Figs. 7-8** and show the resulting hyper-parameter trajectories for (a) AR-d (sparse); (b) Subspace; (c) DNN (MLP) when tuned via SGD. Panels (d,e) parallel those in **Fig. 2a** for (d) $v(t)$ and (e) $u(t)$ outputs of the Izhikevic model. Unlike **Fig. 2a** where we combined across all output channels, we here kept them separate due to the distinct form and role of their dynamics.



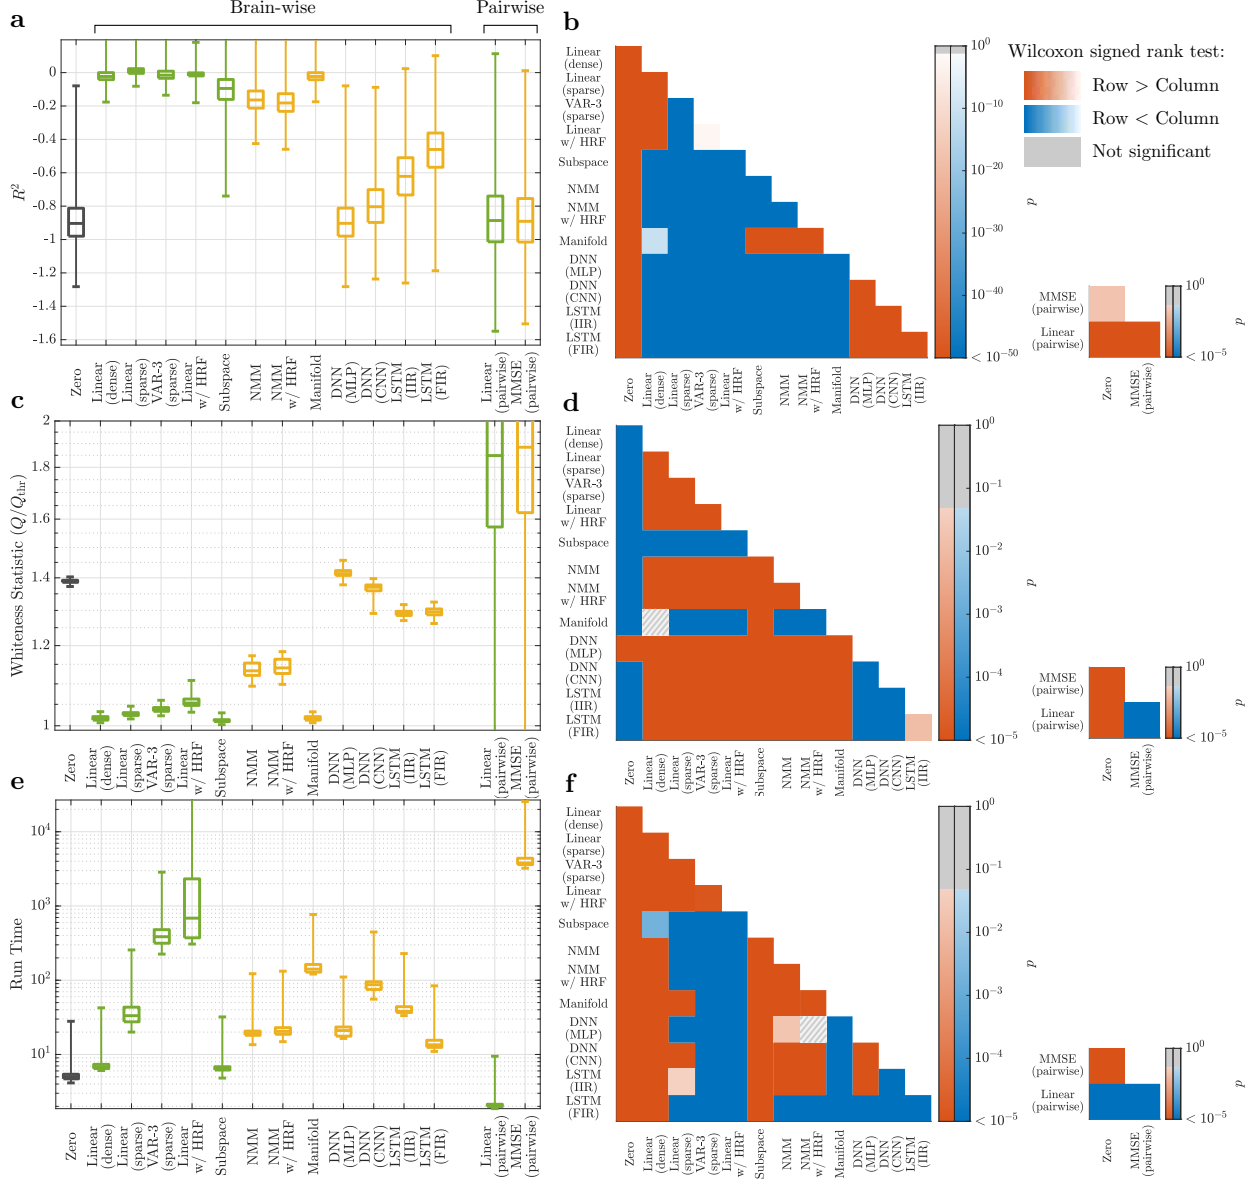
Supplementary Fig. 10 | Comparison of ‘Linear (dense)’ and ‘DNN (MLP)’ models on the logistic map dynamical system. System dynamics are given by $y(t+1) = ry(t)[1 - y(t)]$, $r = 3.7707$. **(a)** Box plots of cross-validated one-step-ahead prediction R^2 . The neural network model achieves near perfect R^2 even with 10 hidden units while the linear model achieves $R^2 \simeq 0.5$. **(b)** The approximations that each model provides to the nonlinear function $y \mapsto ry(1 - y)$. **(c)** The cross-validated one-step-ahead predictions of each model for the first 50 samples of a random run.

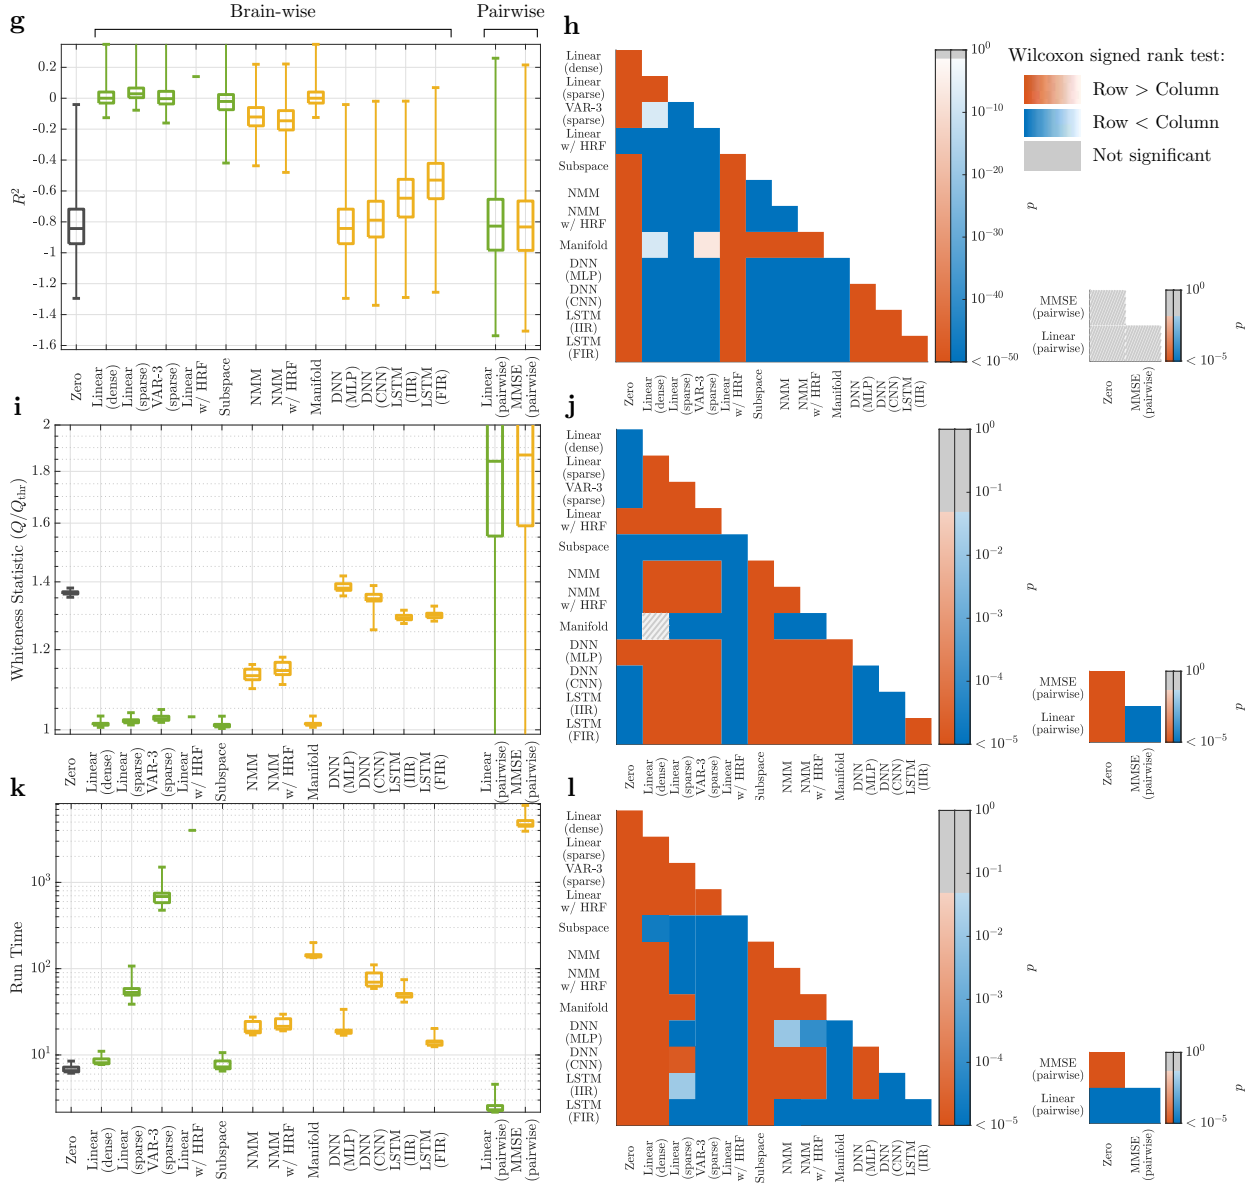


Supplementary Fig. 11 | Comparing ‘DNN (MLP)’ model with a linear activation function against ‘Linear (dense)’ and ‘Linear (sparse)’. (a) The comparison results for the ‘DNN (MLP)’ model with a linear activation function and hidden depth of 0. The network therefore only consists of input and output layers and a fully connected layer in between. The right panel displays the same information as the left panel except for using violin plots. Red crosses and green squares show means and medians, respectively. (b) Hyper-parameter tuning for a ‘DNN (MLP)’ model with a linear activation function and hidden depth of 1. W denotes the width of the hidden layer. As with hyper-parameter tunings in the main text, 100 subjects whose data is not used in subsequent model comparisons are used for hyper-parameter tuning to avoid potential over-fitting to hyper-parameters. (c) Similar to (a) for but for a ‘DNN (MLP)’ model with a hidden depth of 1.

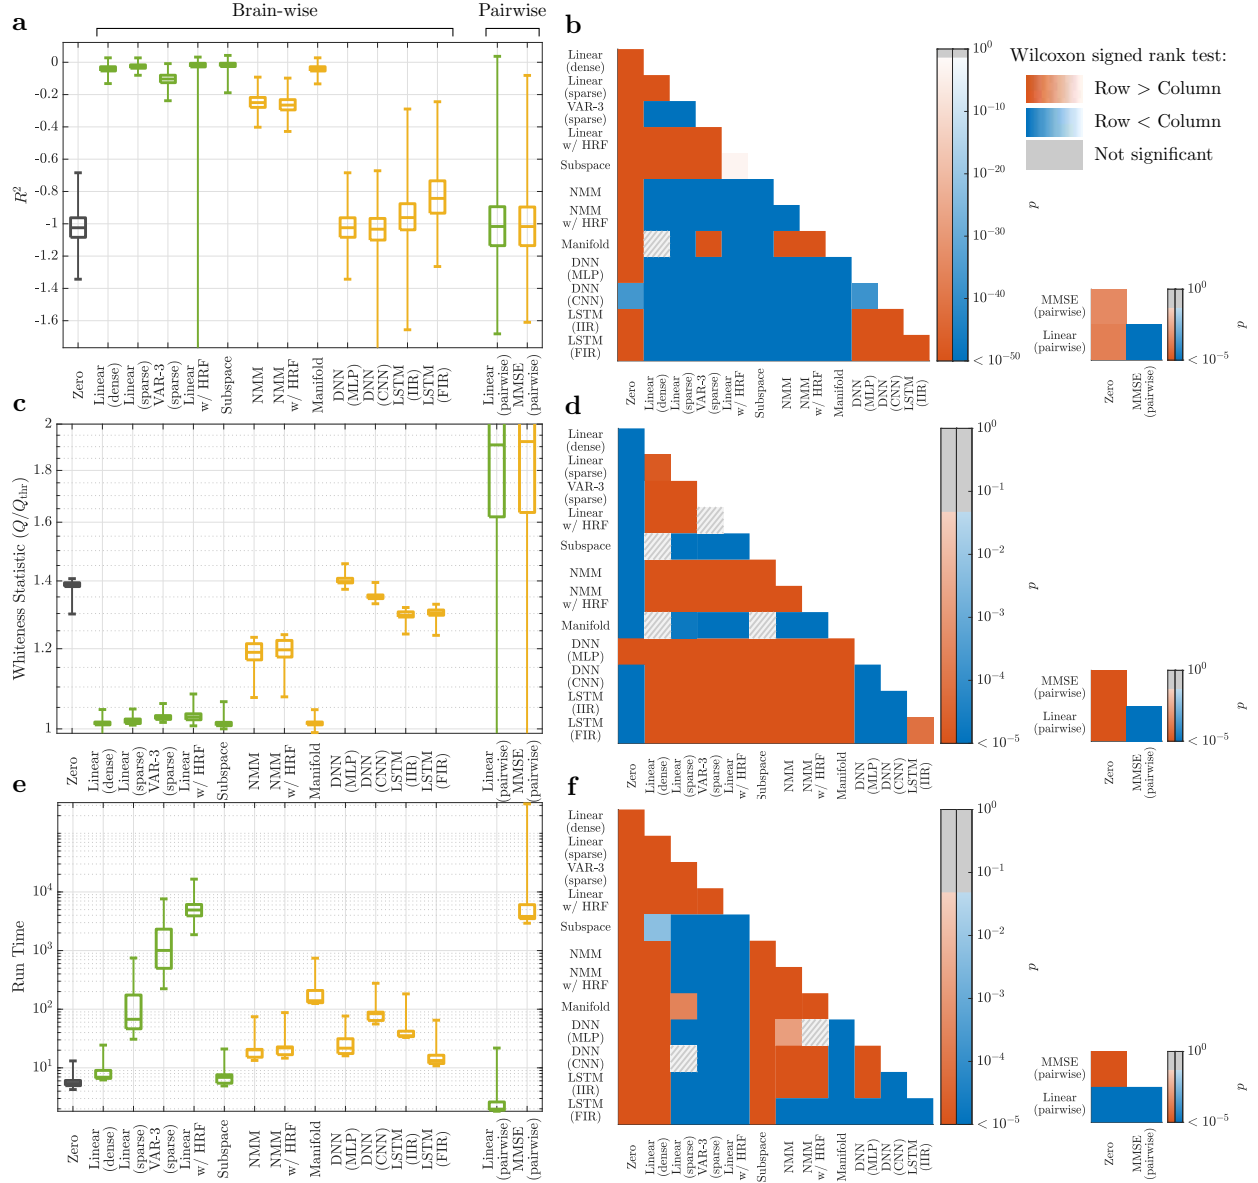


Supplementary Fig. 12 | Linear vs. nonlinear models of finely-parcellated rsfMRI activity. 400 cortical parcels (Schaefer 400x17 [4]) and 50 subcortical ones (Melbourne Scale III [5]) were used. Panels and details parallel those in **Fig. 2** in the main text, except that only data from 32 randomly selected subjects and a single-fold cross-validation has been used to reduce computational complexity. The half-session used as the test for each subject has been selected at random and the remaining 7 half-sessions have been used for training, as in the main text. The model with the highest R^2 is now the ‘Linear (sparse)’ even though ‘VAR-3 (sparse)’ still has whiter residuals.

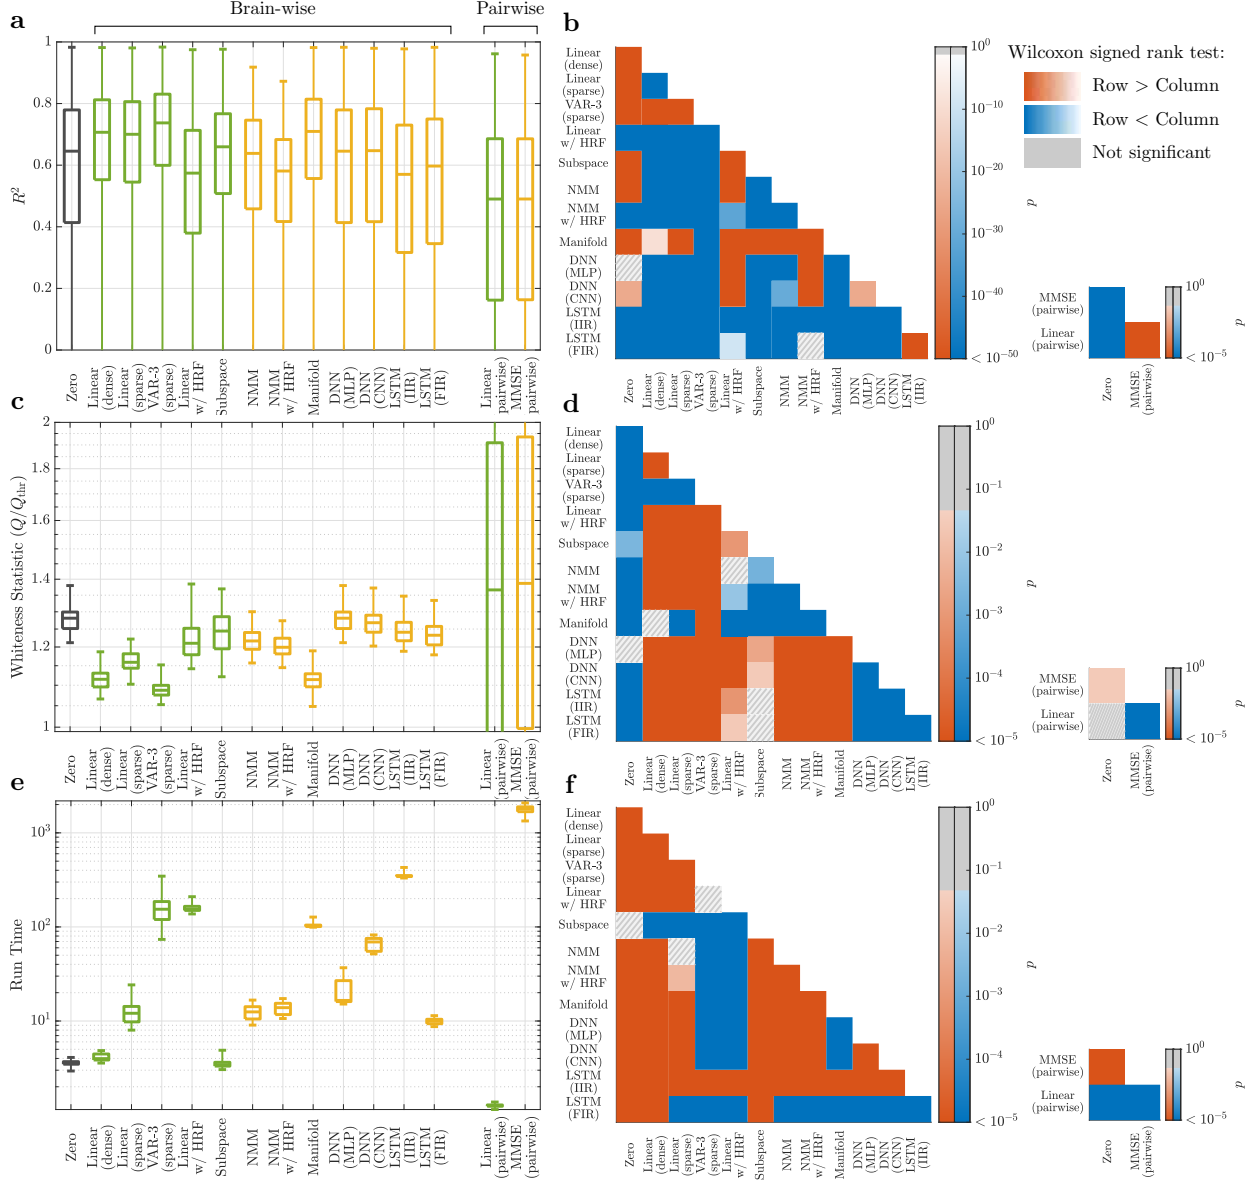




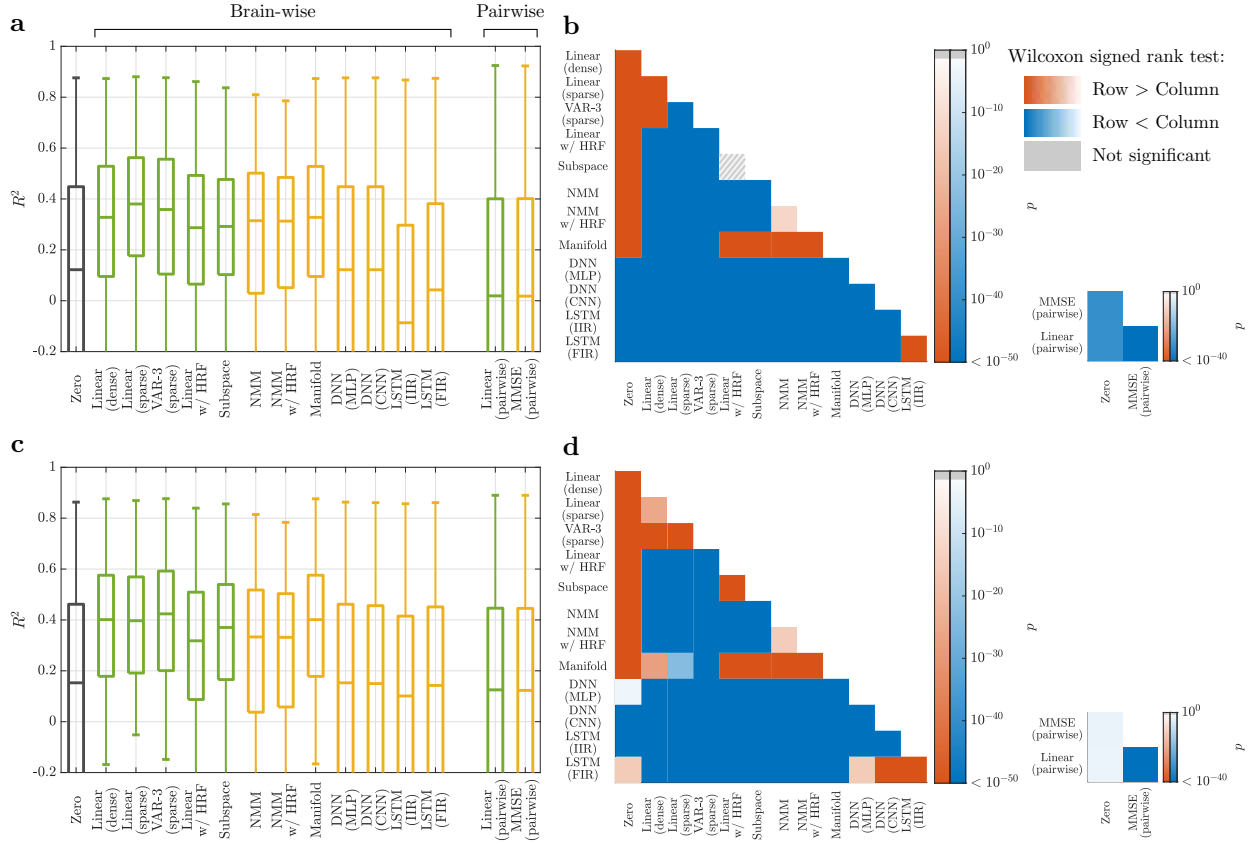
Supplementary Fig. 13 | Linear vs. nonlinear models of unparcellated cortical rsfMRI activity. To reduce computational complexity and be able to fit and validate all model families, we considered only vertices taken from two randomly selected cortical parcels: **(a-f)** a left somato-motor area ‘17Networks_LH.SomMotB_S2’ consisting of 152 vertices, and **(g-l)** a right precuneus/posterior cingulate cortex area ‘17Networks_RH.DefaultA_pCunPCC_1’ consisting of 177 vertices [4]). Panels and details parallel those in **Supplementary Fig. 12**.



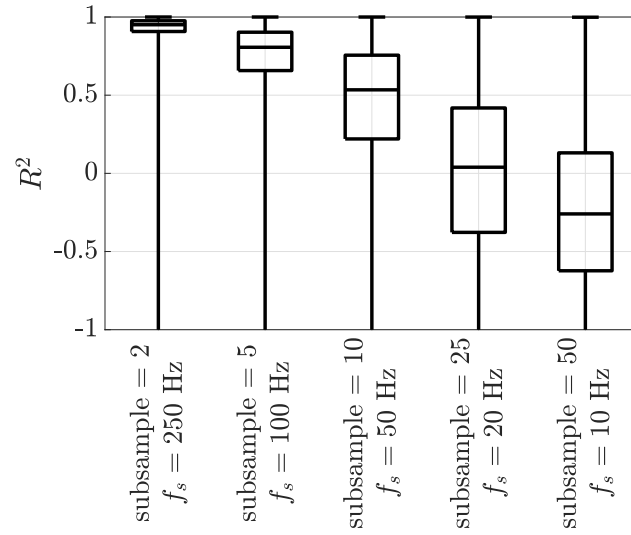
Supplementary Fig. 14 | Linear vs. nonlinear models of unparcellated subcortical rsfMRI activity. To reduce computational complexity and be able to fit and validate all model families, we considered only voxels from one randomly selected subcortical parcel (left dorsoanterior caudate, ‘CAU-DA-lh’ [5], consisting of 154 voxels). Panels and details parallel those in **Supplementary Fig. 12**.



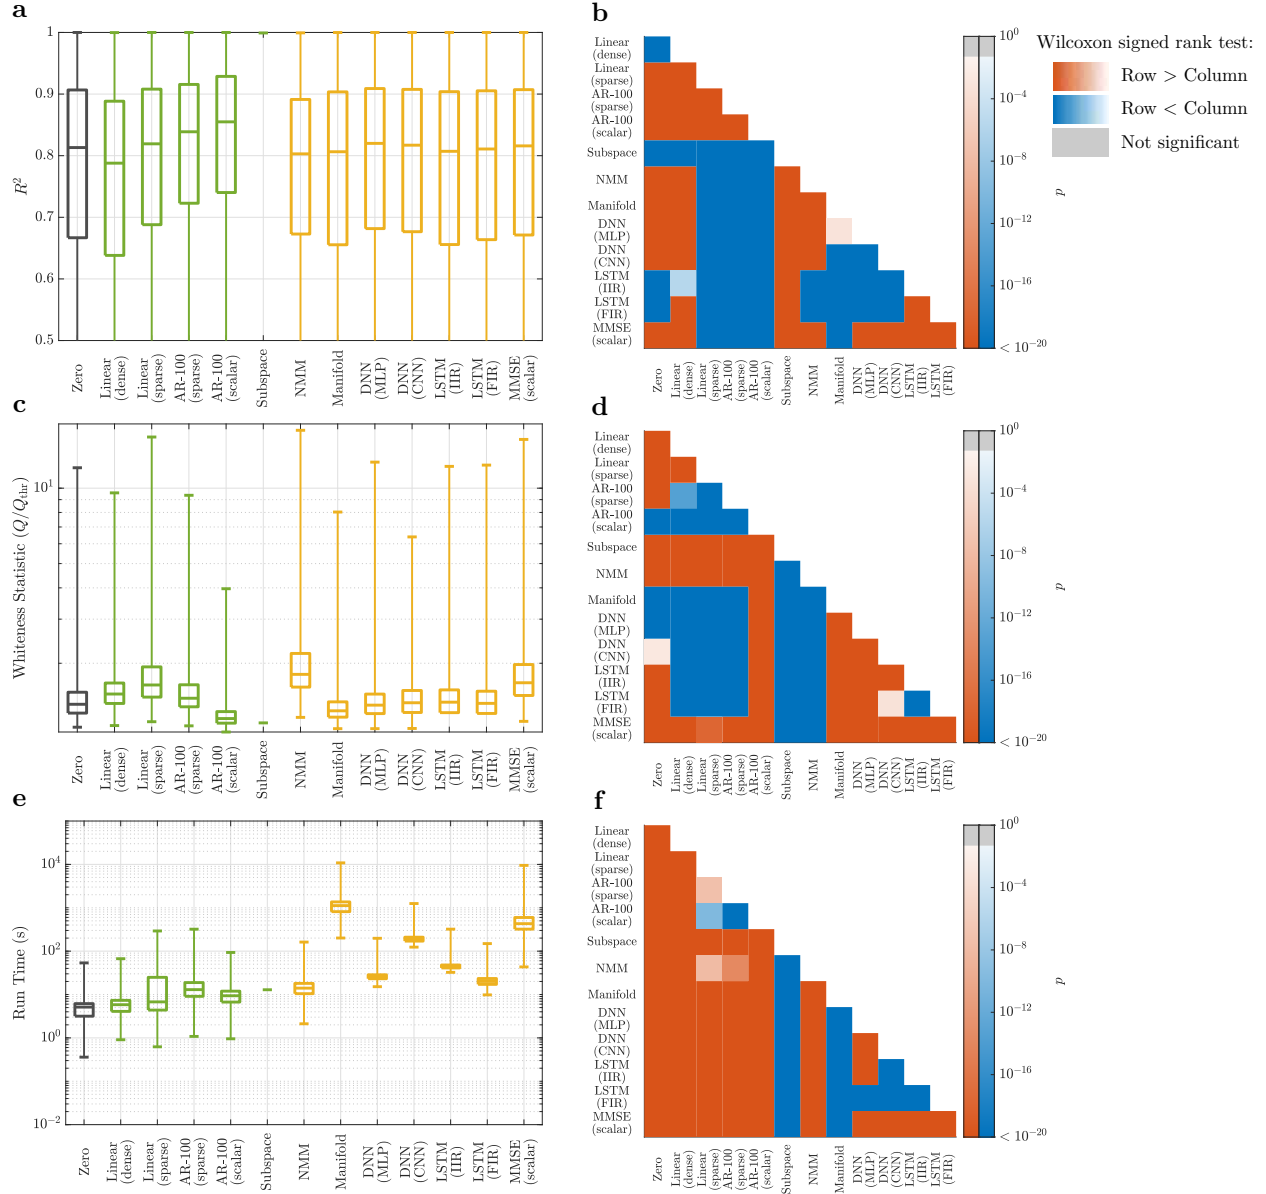
Supplementary Fig. 15 | Linear vs. nonlinear models of minimally pre-processed rsfMRI activity. Data from HCP minimally preprocessed data without ICA-FIX denoising has been used in order to ensure that the observed linearity has not stemmed from ICA-FIX. Panels parallel those in **Fig. 2** in the main text, except that data from a random selection of 10% (70) of subjects is used to reduce computational complexity.



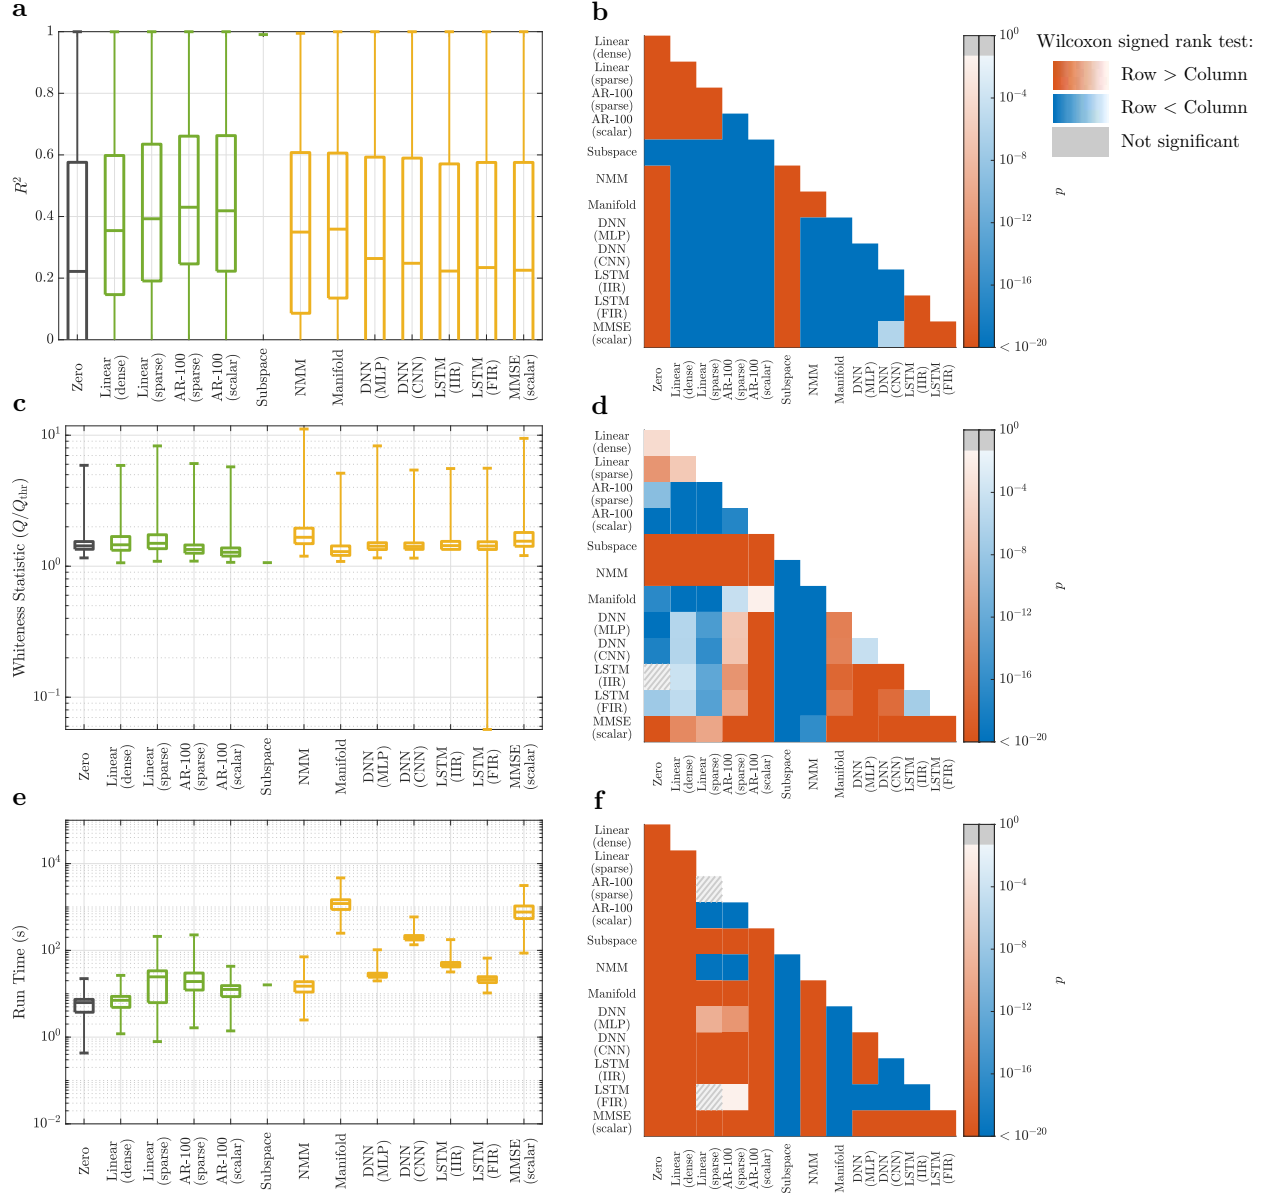
Supplementary Fig. 16 | Test-retest validation of model comparisons for rsfMRI data. Panels in each row parallel those in **Fig. 2a,b** in the main text. Panels (a,b) show the results of a session-wise approach, where the data from a single resting state session is used for both training (75%) and test (25%). Panels (c,d) show the results of a leave-one-session-out approach where 3 resting state sessions (chosen at random) from a subject are used for training and the last session from the same subject is used for test. Data from a random selection of 10% (70) of subjects is used to reduce computational complexity.



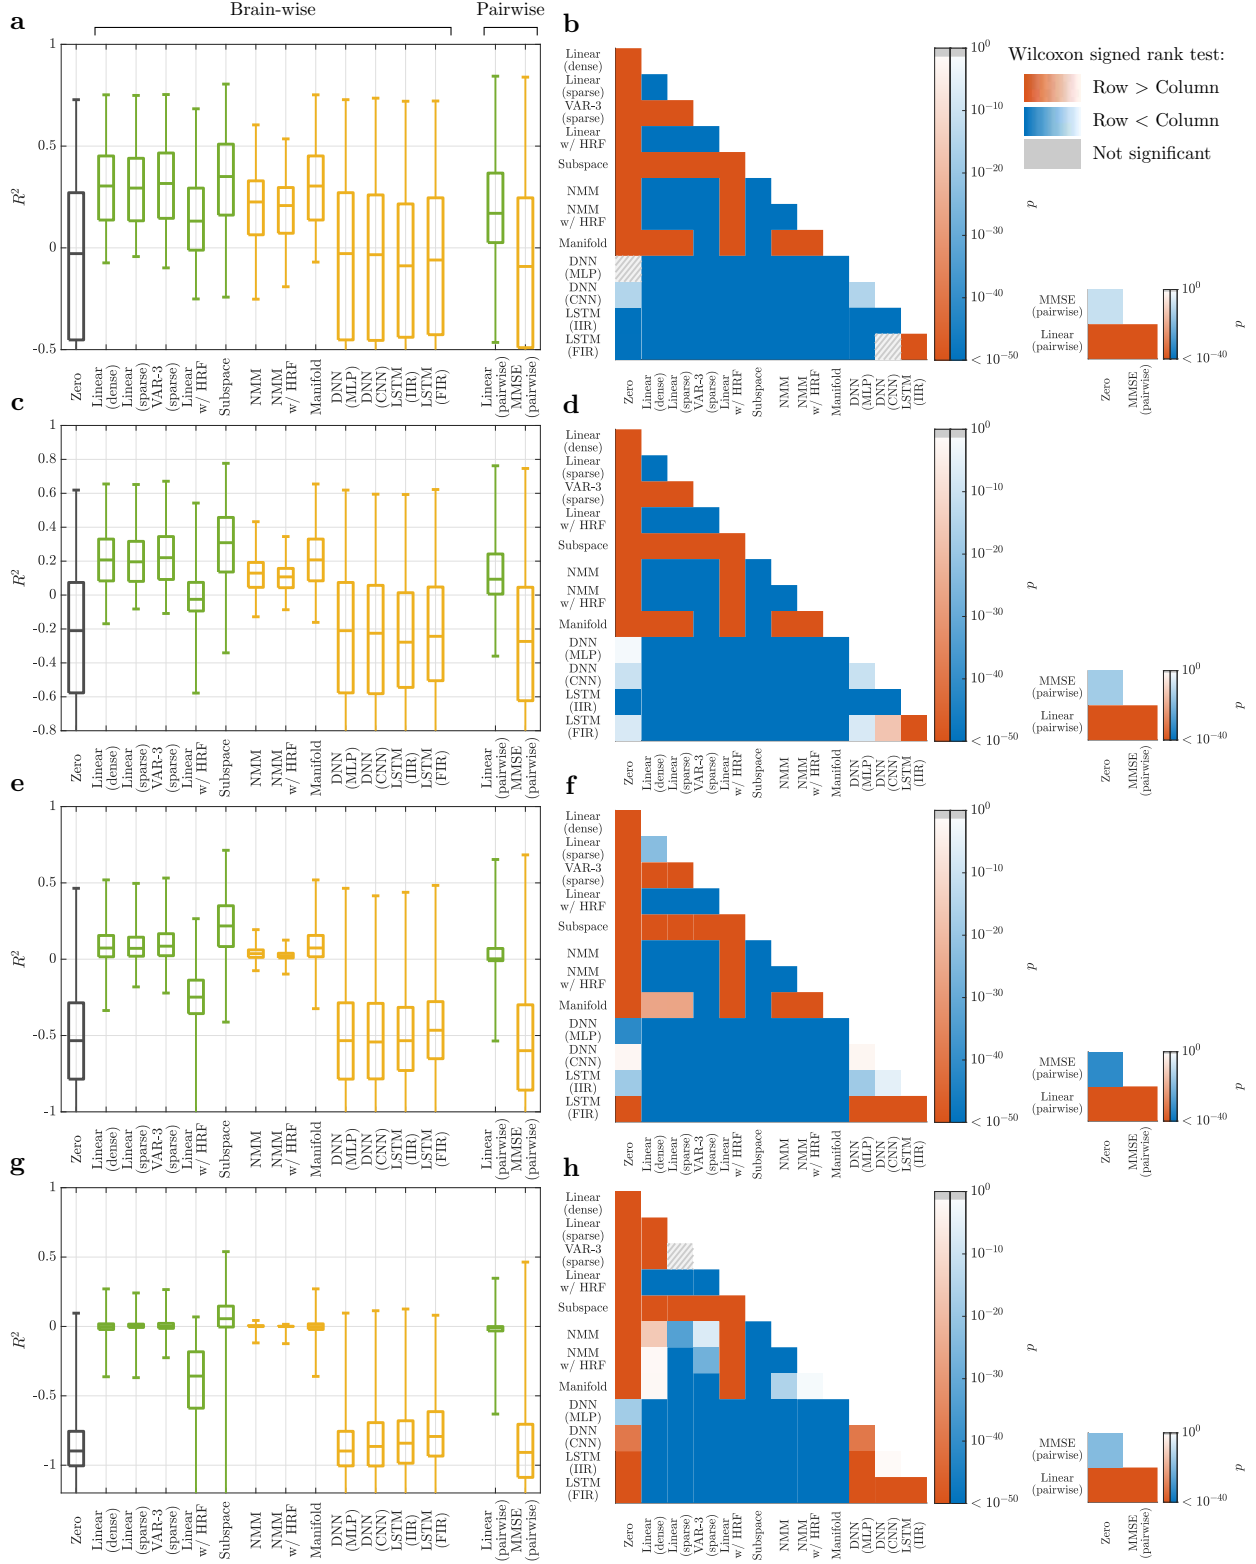
Supplementary Fig. 17 | The channel-wise R^2 distribution of the zero model for iEEG data with different subsampling ratios and the corresponding sampling frequency. As expected, higher subsampling results in less smooth time series, which in turn results in lower ‘Zero’ R^2 , but also allows for using data from longer time intervals in model fitting and validation with the same amount of memory.



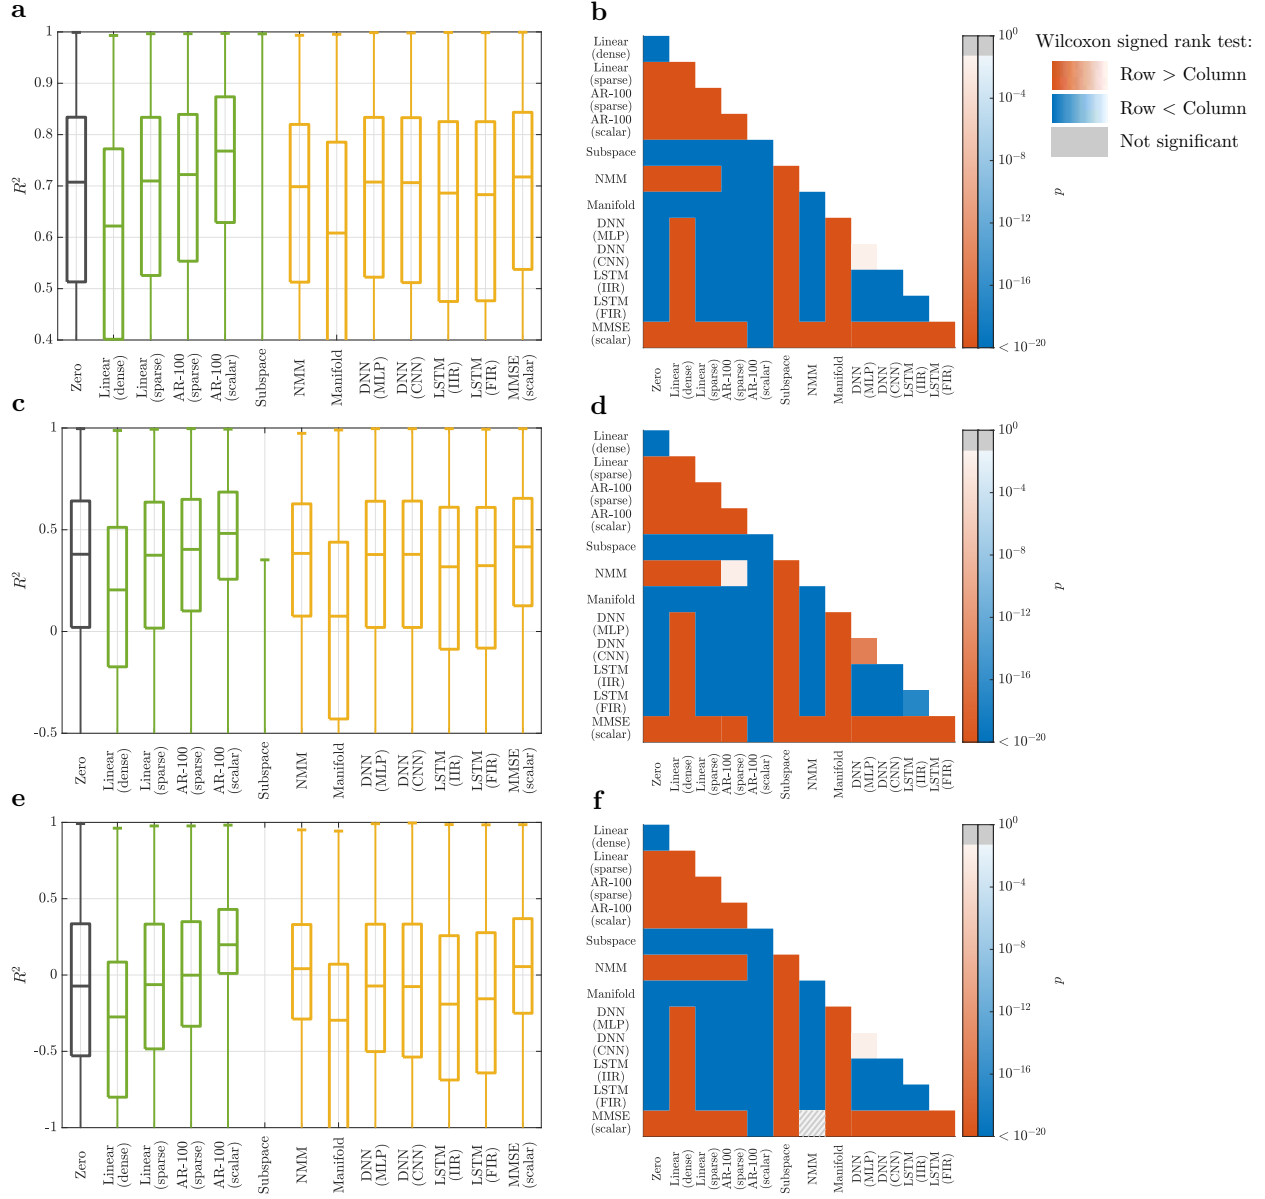
Supplementary Fig. 18 | Linear versus nonlinear models of 5-fold subsampled rsEEG activity. Panels parallel those in **Fig. 3** in the main text. In panel (a), the box plot for the subspace method only has the top line (100th percentile) because for more than 75% of data segments the subspace method was unable to complete (either hung indefinitely or caused MATLAB to crash). For such cases we assign $R^2 = -\infty$, $Q = +\infty$, run time = $+\infty$ for the subspace method, causing its boxplot to miss the third quartile and anything below that. A similar situation holds in panels (c) and (e).



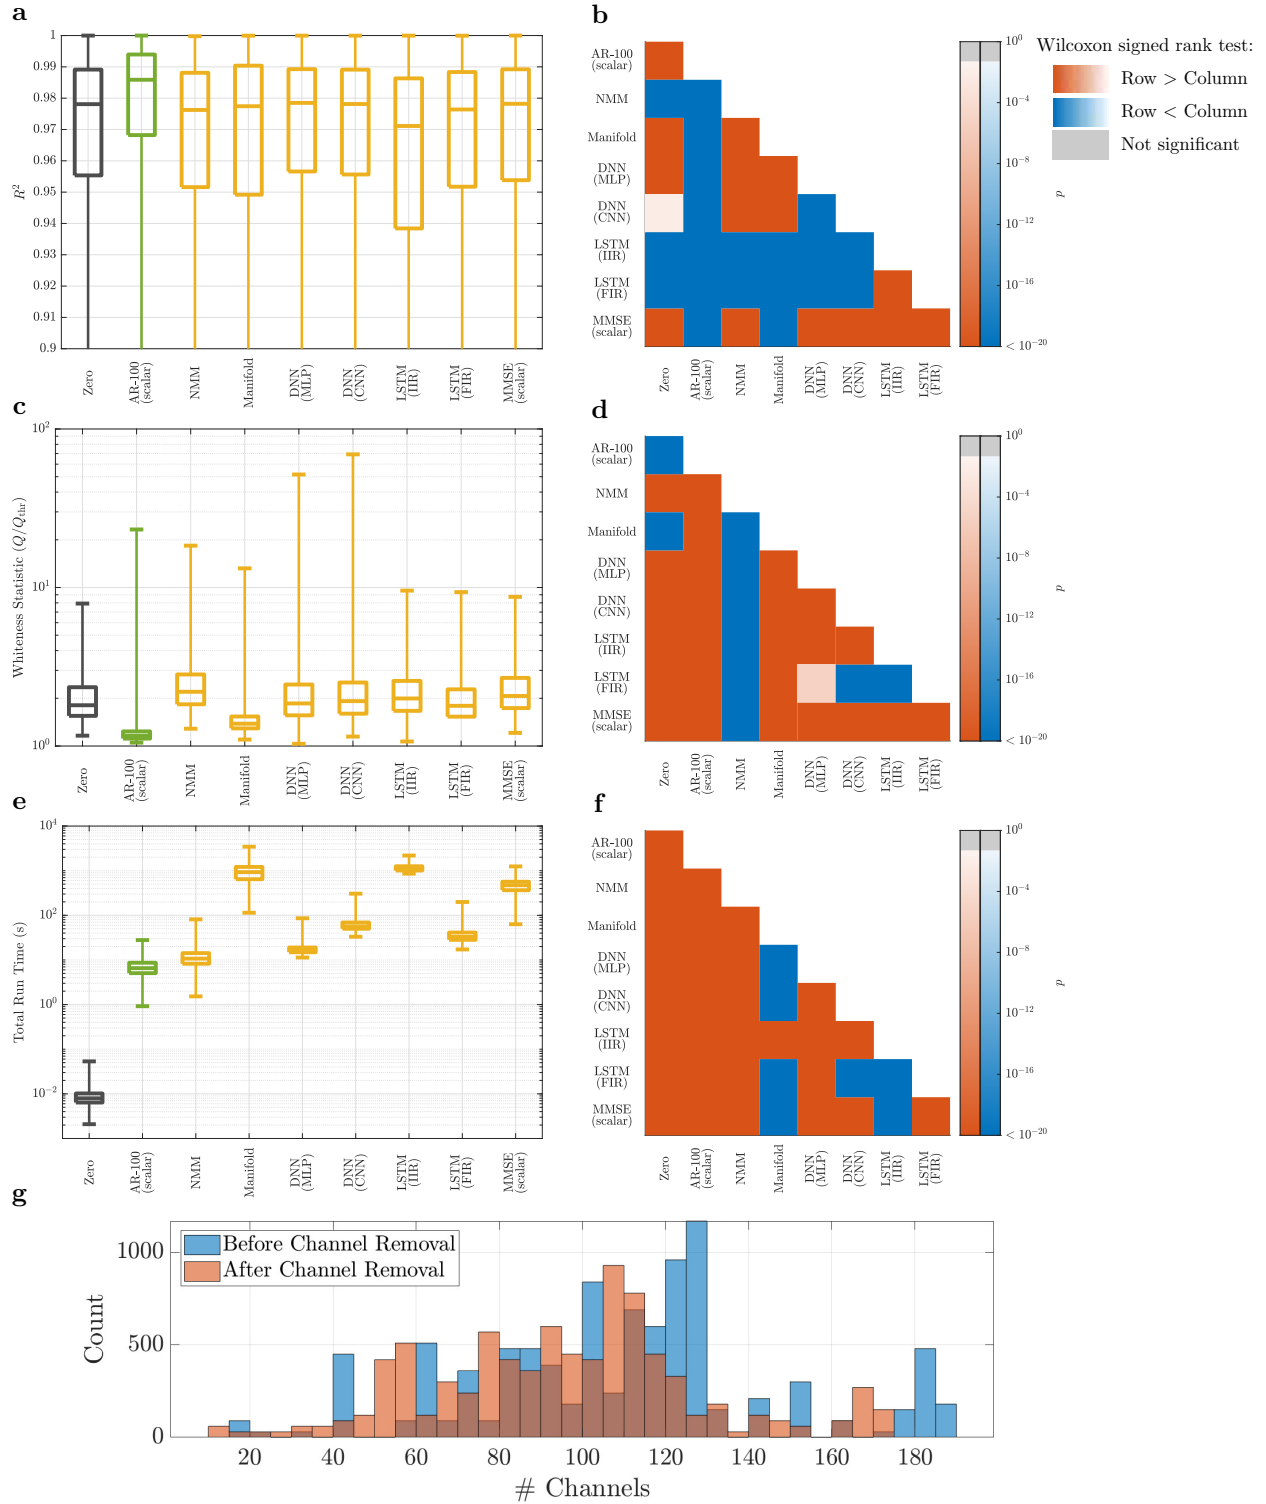
Supplementary Fig. 19 | Linear versus nonlinear models of 25-fold subsampled rsfEEG activity. Panels and details parallel those in **Supplementary Fig. 18**. Note that here the ‘AR-100 (sparse)’ (which includes network interactions) has the highest R^2 distribution, even though the ‘AR-100 (scalar)’ model still has the whitest residuals.



Supplementary Fig. 20 | k -step ahead prediction of rsfMRI data. Panels in each row parallel those in **Fig. 2a,b** in the main text for (a,b) $k = 2$, (c,d) $k = 3$, (e,f) $k = 5$, (g,h) $k = 10$. Data from a random selection of 10% (70) of the subjects is used to reduce computational complexity.



Supplementary Fig. 21 | k -step ahead prediction of rsEEG data. Panels in each row parallel those in Fig. 3a,b in the main text for (a,b) $k = 5$, (c,d) $k = 10$, (e,f) $k = 20$. Data from a random selection of 84 iEEG segments (1% of the total data used in the main text) is used to reduce computational complexity.



Supplementary Fig. 22 | Model comparisons on iEEG data without channel removal. Panels (a-f) parallel those in **Fig. 3** in the main text, except that here no channels are removed due to being noisy as done in the main text and explained under Methods and we only compared the top-performing linear model (‘AR-100 (scalar)’) with nonlinear models. The result echos the results obtained in the main text and ensures that channel removal (although performed according to standard practices and highly advisable from a data quality perspective) had not confounded the superiority of the linear models. (g) The distributions of channel counts across all subjects before and after the removal of noisy channels. The two distributions have medians 110.5 and 98 and means 108.8 and 95.4, respectively.

Supplementary References

- [1] Ciric, R. *et al.* Mitigating head motion artifact in functional connectivity mri. *Nature protocols* **13**, 2801–2826 (2018).
- [2] Oppenheim, A. V., Willsky, A. S. & Nawab, S. H. *Signals and Systems* Prentice-Hall signal processing series (Prentice Hall, 1997). URL <https://books.google.com/books?id=LwQqAQAAMAAJ>.
- [3] Ljung, L. System identification: theory for the user. *PTR Prentice Hall, Upper Saddle River, NJ* 1–14 (1999).
- [4] Schaefer, A. *et al.* Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebral Cortex* **28**, 3095–3114 (2018).
- [5] Tian, Y., Margulies, D. S., Breakspear, M. & Zalesky, A. Hierarchical organization of the human subcortex unveiled with functional connectivity gradients. *bioRxiv* (2020).