Linearizing and forecasting: a reservoir computing route to digital twins of the brain

Gabriele Di Antonio[1,2,3], Tommaso Gili[4], Andrea Gabrielli[1,2], and Maurizio Mattia[3,*]

[1] *"Enrico Fermi" Research Center - CREF, 00184 Rome, Italy*
[2] *Dip. di Ingegneria Civile, Informatica e delle Tecnologie Aeronautiche, Università degli Studi "Roma Tre", 00146 Rome, Italy*
[3] *Natl. Center for Radiation Protection and Computational Physics, Istituto Superiore di Sanità, 00161 Rome, Italy*
[4] *Networks Unit, IMT Scuola Alti Studi Lucca, 55100 Lucca, Italy*
[*] *Corresponding author: maurizio.mattia@iss.it*

October 2, 2025

Abstract

Exploring the dynamics of a complex system, such as the human brain, poses significant challenges due to inherent uncertainties and limited data. In this study, we enhance the capabilities of noisy linear recurrent neural networks (lRNNs) within the reservoir computing framework, demonstrating their effectiveness in creating autonomous *in silico* replicas – digital-twins – of brain activity. Our findings reveal that the poles of the Laplace transform of high-dimensional inferred lRNNs are directly linked to the spectral properties of observed systems and to the kernels of auto-regressive models. Applying this theoretical framework to resting-state fMRI, we successfully predict and decompose BOLD signals into spatiotemporal modes of a low-dimensional latent state space confined around a single equilibrium point. lRNNs provide an interpretable proxy for clustering among subjects and different brain areas. This adaptable digital-twin framework not only enables virtual experiments but also offers computational efficiency for real-time learning, highlighting its potential for personalized medicine and intervention strategies.

Introduction

The study of complex dynamical systems has a central role in contemporary scientific research. Over recent decades, a plethora of data-driven methodologies has emerged, denoting a field in rapid evolution that is capable to addressing systems with increasing complexity. Among these methodologies, machine learning techniques have become increasingly popular due to their effectiveness in modeling complex datasets and deliver accurate predictions (1, 2, 3).

Within the spectrum of machine learning techniques, reservoir computing (RC) has distinguished itself as a powerful approach for processing temporal data. RC leverages high-dimensional recurrent neural networks (RNNs) with fixed internal couplings to transform input time series into rich, high-dimensional representations (4, 5). This transformation enables the prediction of future samples by simply reading out the current state of the network, effectively capturing the system's dynamics (6, 7, 8). The elegance of this approach lies in conceptualizing the readout as a projection onto a manifold learned through linear regression, echoing the efficient coding strategies observed in biological neural networks (9, 10, 11). Consequently, RNNs have become invaluable tools in neuroscience, aiding in hypothesis generation and providing analytical frameworks for understanding neural computations (12, 13, 14, 15).

However, as machine learning has become increasingly popular, some clarity of understanding may have been lost in the pursuit of better predictions. While the inferred models are often quite accurate, they can sometimes lack interpretability (16, 17, 18). A valuable framework to address this limitation is provided by Koopman theory, which offers clear insights into the system dynamics by mapping nonlinear systems into simple yet higher-dimensional dynamics (19, 20, 21, 22). This approach involves the development of data-driven methods aimed at computing finite-dimensional linear representations of nonlinear systems within a suited functional space. Once an equivalent

linear system is inferred, interpretability becomes straightforward, allowing to gain understanding of the time scales and dynamical modes of the observed systems.

The capability to predict future states of a dynamical system based on past observations allows, in principle, to build a replica of the same system. Indeed, predicted observations can be fed back as input resulting in a generative model. Both Koopman-based methods (Fig. 1A) and RC approaches (Fig. 1B) have the potential to generate these digital copies by producing time series statistically equivalent to those generated by the original systems. The inferred generative models can then serve as 'digital twins' – a concept that has emerged as a transformative paradigm in the study of dynamical and natural systems (23, 24). Such digital twins can be effective tools for analyzing and simulating their physical counterparts. Successful examples include digital copies of brain activity designed to assist neurosurgeons in dealing with drug-resistant epileptic patients (25, 26, 27). Furthermore, these tools can provide a complementary workbench for designing stimulation approaches to be tested *in silico*, ultimately leading to control systems for neurorehabilitation (28).

In this study, we investigate the potential of RNNs with linear units to implement effective digital twins of single-subject brain dynamics, measured as the vascular response to neuronal activity (i.e., the BOLD signal) extracted from functional magnetic resonance imaging (fMRI). We demonstrate the autoregressive nature of this approach and its effectiveness in representing whole-brain activity in a low-dimensional latent state space. Such a reduced representational complexity mitigates the risk of overfitting often associated with one-to-one fitted physical models, especially in realistic scenarios where only a limited number of observations are available and only a partial view of the system can be experimentally accessed. The integration of endogenous memoryless fluctuations is another essential component that enables the digital twin to self-sustain a noisy dynamics statistically equivalent to the one experimentally measured.

Results

Linear RNNs can copy nonlinear systems

We start investigating the reservoir computing abilities of a RNN with N linear units (lRNN). The state of the lRNN is the vector $\boldsymbol{r}(t)$ whose elements are the activities $r_i(t)$ of each unit, following the linear dynamics

$$\tau \dot{\boldsymbol{r}}(t) = \boldsymbol{W}\boldsymbol{r}(t) + \boldsymbol{W}^{\text{in}}\boldsymbol{\omega}(t) - \boldsymbol{r}(t)\,. \tag{1}$$

The synaptic matrix $\boldsymbol{W} \in \mathbb{R}^{N\times N}$ determines the recurrent input $\boldsymbol{W}\boldsymbol{r}$, which, together with the contributions from the external environment $\boldsymbol{W}^{\text{in}}\boldsymbol{\omega}(t)$, composes the total input $\boldsymbol{h}$ that each unit receives. In general, these currents are nonlinearly transformed by an input-output gain function $\Phi(\boldsymbol{h})$ we assume here to be $\Phi(\boldsymbol{h}) = \boldsymbol{h}$ – a condition usually associated with unit activities perturbatively fluctuating around the quiescent state (see Materials and Methods). Network units receive a source of external input due to the M observables $\omega_i(t)$ of the systems under investigation, mediated by the synaptic matrix $\boldsymbol{W}^{\text{in}} \in \mathbb{R}^{N\times M}$. From Eq. (1) the network state at any time t results then to be

$$\boldsymbol{r}(t) = \int_{-\infty}^{t} e^{(\boldsymbol{W}-\boldsymbol{I})(t-t')/\tau}\boldsymbol{W}^{\text{in}}\boldsymbol{\omega}(t')dt'\,, \tag{2}$$

where $\boldsymbol{I} \in \mathbb{R}^{N\times N}$ is the identity matrix.

If the lRNN is capable to display the so-called 'generalized synchronization' property in response to the input $\boldsymbol{W}^{\text{in}}\boldsymbol{\omega}$ (29, 30), $\boldsymbol{r}(t)$ effectively embeds the dynamics of the 'whole' observed system into its state space. This because the lRNN incorporates the 'echo' (31, 4, 32) of past observations implementing a dimensional embedding *a la* Takens (33, 34, 35). In this RC framework, for sufficiently high N a simple linear transformation capable to approximately map the network state to a target output exists

$$\boldsymbol{\omega}(t) \approx \boldsymbol{W}^{\text{out}}\boldsymbol{r}(t) = \int_{0}^{\infty} \boldsymbol{G}(s)\boldsymbol{\omega}(t-s)ds \tag{3}$$

with synaptic weights $\boldsymbol{W}^{\text{out}} \in \mathbb{R}^{M\times N}$ computed via a simple linear regression (4, 5). The linear kernels $\boldsymbol{G}(s) = \boldsymbol{W}^{\text{out}}E^{(\boldsymbol{W}-\boldsymbol{I})(s)/\tau}\boldsymbol{W}^{\text{in}} \in \mathbb{R}^{M\times M}$ (i.e., the Green functions) resulting from Eq. (2) reveals the autoregressive nature of the RC approach in this lRNN case. This concept can be better clarified by resorting to the eigendecomposition of the synaptic matrix $\boldsymbol{W} = \boldsymbol{Q}\boldsymbol{\Lambda}\boldsymbol{Q}^{-1}$. Here $\boldsymbol{\Lambda}$ is the diagonal matrix of the eigenvalues

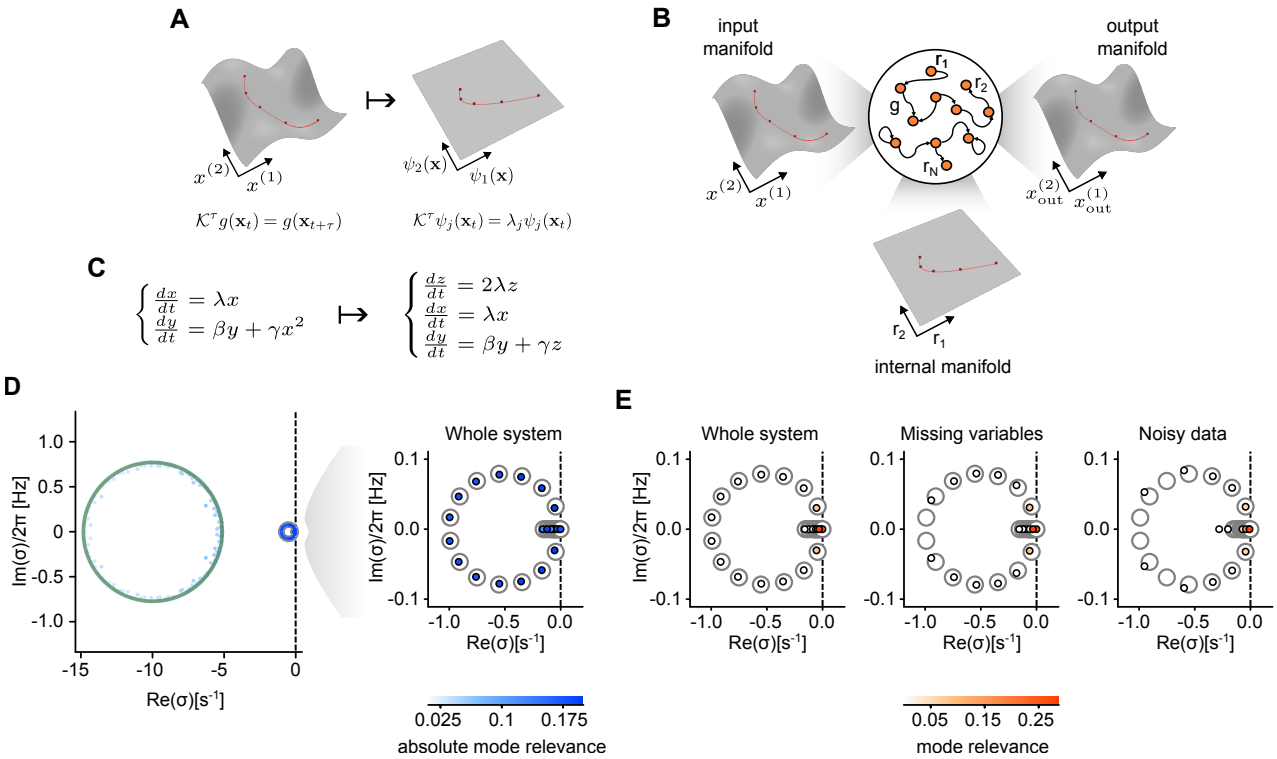

Figure 1: **IRNN encodes a linear representation of the input system**. (**A**) Nonlinear dynamics can be projected onto linear manifolds. Within the infinite functional space of the system observables, the eigenfunctions of the Koopman operator establish a basis for decomposing nonlinear dynamics into simpler linear processes. (**B**) Reservoir computing (RC) with recurrent neural networks (RNNs). After 'learning', the readout weights $\boldsymbol{W}^{\text{out}}$ are set to reproduce the system observables provided as input. The resulting network is composed of linear units (IRNN) and its post-learning autonomous dynamics encapsulates a linear representation of the original system. (**C**) Example of a 2-dimensional nonlinear system that can be decomposed into a set of 3 linearly evolving variables. (**D**) Spectrum of a IRNN with $N = 500$ units trained to mimic a 15-dimensional version of the system in (C) composed of 5 linear x variables and 10 nonlinear y variables. Green circles, spectrum of the IRNN before training. Gray circles, theoretically-derived eigenvalues of the 20-dimensional linearized system to be replicated (see text for details). Small circles, Laplace transform poles from the post-learning autonomous IRNN replicating the example system. Filling colors code the absolute mode relevance $\langle |\Xi_{jk}| \rangle_j$. Right panel, zoom in of the left panel showing the 20 most significant modes. (**E**) Same spectrum as in (D) with filling colors coding the mode relevance $\langle |\Xi_{jk} v_k(0)| \rangle_j$ (left panel, see text for details). Central panel, spectrum of a IRNN trained on incomplete information (i.e., omitting all x and 5 y variables). Right panel, spectrum of a IRNN trained on the same observables but perturbed by noise.

($\Lambda_{ii} = \lambda_i$) and the eigenvectors are the columns of $\boldsymbol{Q}$: $\boldsymbol{W} |Q_i\rangle = \lambda_i |Q_i\rangle$. By applying this decomposition to the kernel expression, a superposition of N exponential modes results:

$$\boldsymbol{G}(s) = \sum_{n=1}^{N} \boldsymbol{J}_n e^{(\lambda_n - 1)s/\tau} \,, \tag{4}$$

where $\boldsymbol{J}_n = \boldsymbol{W}^{\text{out}} |Q_n\rangle \langle Q_n^{-1}| \boldsymbol{W}^{\text{in}}$ are N rank-1 matrices with $\langle Q_n^{-1}|$ being the n-th row of the inverse matrix $\boldsymbol{Q}^{-1}$ (see Materials and Methods). The convolution in Eq. (3) is integrable if the spectral radius $\rho = \max_{n\in[1,N]} \text{Re}\, \lambda_n$ of $\boldsymbol{W}$ is smaller than 1. Under this condition, unperturbed IRNNs ($\boldsymbol{\omega} = \boldsymbol{0}$) have a stable equilibrium point at $\boldsymbol{r} = \boldsymbol{0}$, and self-consistency equation (3) describes a continuous-time autoregressive (AR) model with order N (number of exponential modes in $\boldsymbol{G}(s)$). This relationship between RC with IRNNs and AR models generalizes previous results obtained in the discrete-time domain (36, 37). As we will see later, the order of the equivalent AR model

can be much smaller than N as many $\boldsymbol{J}_n$ are usually close to $\mathbf{0}$.

By replacing the observables $\boldsymbol{\omega}(t)$ provided as input with the reconstructed once from Eq. (3), the IRNN results to follow the autonomous dynamics

$$\tau \dot{\boldsymbol{r}} = \tilde{\boldsymbol{W}} \boldsymbol{r} - \boldsymbol{r} \tag{5}$$

where $\tilde{\boldsymbol{W}} = \boldsymbol{W} + \boldsymbol{W}^{\text{in}} \boldsymbol{W}^{\text{out}}$ is the updated (i.e., 'learned') synaptic matrix. If the network can correctly predict the future of the target output, Eq. (5) provides a linear representation of the system under analysis. Relying as above on the eigenmode decomposition $\tilde{\boldsymbol{W}} = \tilde{\boldsymbol{Q}} \tilde{\boldsymbol{\Lambda}} \tilde{\boldsymbol{Q}}^{-1}$, the projections $\boldsymbol{v}(t) \equiv \tilde{\boldsymbol{Q}}^{-1} \boldsymbol{r}(t)$ evolve in time as a set of decoupled variables:

$$v_k(t) = v_k(0) e^{(\tilde{\lambda}_k - 1)t/\tau} \quad \forall k \in [1, N]. \tag{6}$$

This last step makes it apparent the relationship with the Koopman-operator theory, as a suited linear combination of these modes eventually implements an effective decomposition of the time series $\boldsymbol{\omega}(t)$:

$$\boldsymbol{\omega}(t) = \boldsymbol{W}^{\text{out}} \boldsymbol{r}(t) \equiv \boldsymbol{\Xi} \boldsymbol{v}(t)\,, \tag{7}$$

with $\boldsymbol{\Xi} = \boldsymbol{W}^{\text{out}} \tilde{\boldsymbol{Q}}$. The Laplace transform of the reconstructed observable will then have N poles $\sigma_k = (\tilde{\lambda}_k - 1)/\tau$ with their related residues, allowing us to represent the dynamical properties of the 'digital twin' (i.e., the autonomous IRNN), and linking it to AR models and finite approximations of the Koopman operator (see Materials and Methods). Indeed, the k-th residue is proportional to $\Xi_{jk} v_k(0)$ and provides in absolute value the relevance of the k-th mode in describing the observable ω_j, given the initial network state $\boldsymbol{r}(0)$.

As a testing ground for the theoretical framework derived above we consider the following nonlinear system:

$$\begin{cases} \dot{\boldsymbol{x}} = \boldsymbol{C}_x \boldsymbol{x} \\ \dot{\boldsymbol{y}} = \boldsymbol{C}_y \boldsymbol{y} + \boldsymbol{C}_{yx} \boldsymbol{x} \odot \boldsymbol{x} \end{cases} \tag{8}$$

Choosing $\boldsymbol{C}_x$ as a diagonal matrix, this system can be mapped onto a linear one by introducing a set of new observables $\boldsymbol{z} = \boldsymbol{x} \odot \boldsymbol{x} = \mathbf{diag}(\boldsymbol{x})\, \boldsymbol{x}$ (Fig. 1C). In this case, the related infinitesimal generator of the Koopman operators is equivalent to the finite matrix $\boldsymbol{K}$ with diagonal blocks $\boldsymbol{C}_x$, $2\boldsymbol{C}_x$, $\boldsymbol{C}_y$.

We trained a IRNN to replicate the relaxation dynamics of this system from the observables $\boldsymbol{\omega} = [\boldsymbol{x}; \boldsymbol{y}]$ with $\boldsymbol{x} \in \mathbb{R}^5$ and $\boldsymbol{y} \in \mathbb{R}^{10}$, starting from a random initial condition $\boldsymbol{\omega}(0)$. As initial synaptic matrix $\boldsymbol{W}$ we chosen the one associated to ring-like connectivity (i.e., $W_{i,i+1} = W_{N,1} = \rho$ are the only non-zero elements) known to have optimal performances in the case of linear units (38). In Fig. 1D the poles σ_k of the reconstructed observables from Eq. (7) are shown together with the related mode relevance (color intensity). Despite the nonlinearity of the system to replicate, the overlap between the poles σ_k and the theoretical spectrum of $\boldsymbol{K}$ is remarkable. Furthermore, only 20 out of the $N = 500$ modes have $\langle |\Xi_{jk}| \rangle_j$ significantly different from 0 (blue-colored circles).

Decomposition modes display damped oscillations as $\tilde{\lambda}_k \in \mathbb{C}$ and a few of them, the slowest ones, contribute significantly to the dynamics (Fig. 1E, red-colored small circles). Such low-dimensionality is not altered even if during learning the IRNN receives noisy-contaminated observables or some of them are not provided as input. Indeed, the autonomous IRNN continues to predict accurately the system's future also under these conditions, although the less relevant poles move away from the imaginary axis (missing match between small and large circles).

IRNNs successfully forecast fMRI time series

To explore the applicability of the RC with IRNN, we examined time series data from a real-world physical system. We considered the brain activity of 20 healthy subjects at rest (i.e., in a state of quiet wakefulness), measured through fMRI. The experimental time series was derived from the BOLD signals recorded from a subset of voxels, which serve as proxies for neuronal activity across 11 cortical areas in the language network (Fig. 2A, see Materials and Methods).

The BOLD time series underwent preprocessing using principal component analysis (PCA) to reduce redundancy and achieve dimensionality reduction, which was found to be relatively high (see Fig. 2B, C). Specifically, we assessed the correlation between the information lost during dimensionality reduction and the actual signals. Our analysis revealed that the first principal components (PCs), which explain 99.5% of the variance (Fig. 2D,E), excluded a residual activity that was indistinguishable from white noise. Based on this criterion, we determined

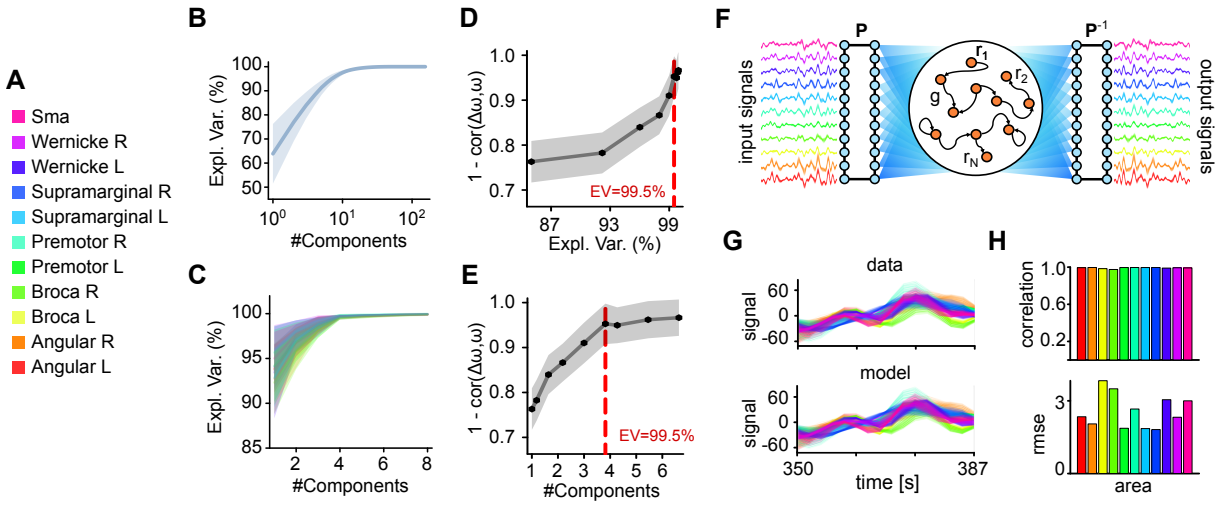

Figure 2: **Compression and prediction of resting state fMRI activity**. (**A**) Color legend associated with the 11 cortical areas taken from the language network. (**B**) Explained variance (EV) from PCA of the fMRI BOLD signals of the voxels in the language network. Solid line, average across the 20 considered subjects. Shaded area, standard deviation across subjects. (**C**) EV of the PCA for each cortical area. Colors as in (A). (**D** and **E**) 1 minus the correlation between the residuals from the PCA reconstruction ($\Delta\omega$) and the original data (ω) per single area in the last (unseen) 37.5 seconds. A value equal to 1 means that the remaining information is just white noise. Vertical dashed lines, explained variance (D) and number of first PCs (E) required to reach the EV of 99.5%. (**F**) Schematic design of the used IRNN. The input from each area is compressed to have an EV of 99.5%. $\boldsymbol{P}$ is the matrix operating this linear transform leading to an average of 44 observables. The IRNN output reconstructing these observable is then uncompressed applying $\boldsymbol{P}^{-1}$. (**G**) Example of forecasted (bottom) and original (top) activities during the test window (last 37.5 seconds). Each curve is the BOLD signal of a voxel with color indicating the cortical area of belonging. (**H**) Forecast performances of the post-learning autonomous IRNN per area. Correlations between predicted and experimental activity (top) and root mean square error (bottom).

that, on average, only four PCs per cortical area should be considered. Consequently, the average dimensionality of the dataset was $M \approx 44$ PCs per subject (refer to Materials and Methods for details). This relatively high-dimensional time series served as the input for the IRNN. The readout synaptic matrix $\boldsymbol{W}^{\text{out}}$ was computed over the first 280 seconds (Fig. 2F).

The decision to utilize a transformation of brain activity that results in information loss may appear as counter-productive. However, projecting onto the first PCs was expected to be less sensitivity to the uncertainties inherent in the measured BOLD signals. By employing these first PCs as inputs, we provided the IRNN with denoised representations of brain activity. Moreover, thanks to the echo state property (31, 4, 32), the missing information regarding the state of the observed system can be effectively recovered through the dimensional embedding that the IRNN performs.

As previously described, with the learned $\boldsymbol{W}^{\text{out}}$ – which varies from subject to subject – we established a closed loop between the input and output of the IRNN, effectively creating an autonomous digital twin of the brain network under investigation. This approach yielded a remarkable overlap between the reconstructed and measured time series (see Fig. 2G for an example subject). At the population level, we evaluated the forecasting performance by calculating the correlations between replicated and experimental BOLD signals, as well as the root mean square errors (RMSEs) for the final unseen 37.5 seconds of the time series (Fig. 2H). From this perspective, the inferred digital twin demonstrated high-quality reconstruction of brain activity, further validating its effectiveness.

IRNN as a digital twin of brain activity

The IRNN dynamics described thus far is dissipative, meaning that the state of the network ultimately relaxes to $\mathbf{0}$. However, brain activity associated to the BOLD signals fluctuates without rest. To reconcile this discrepancy

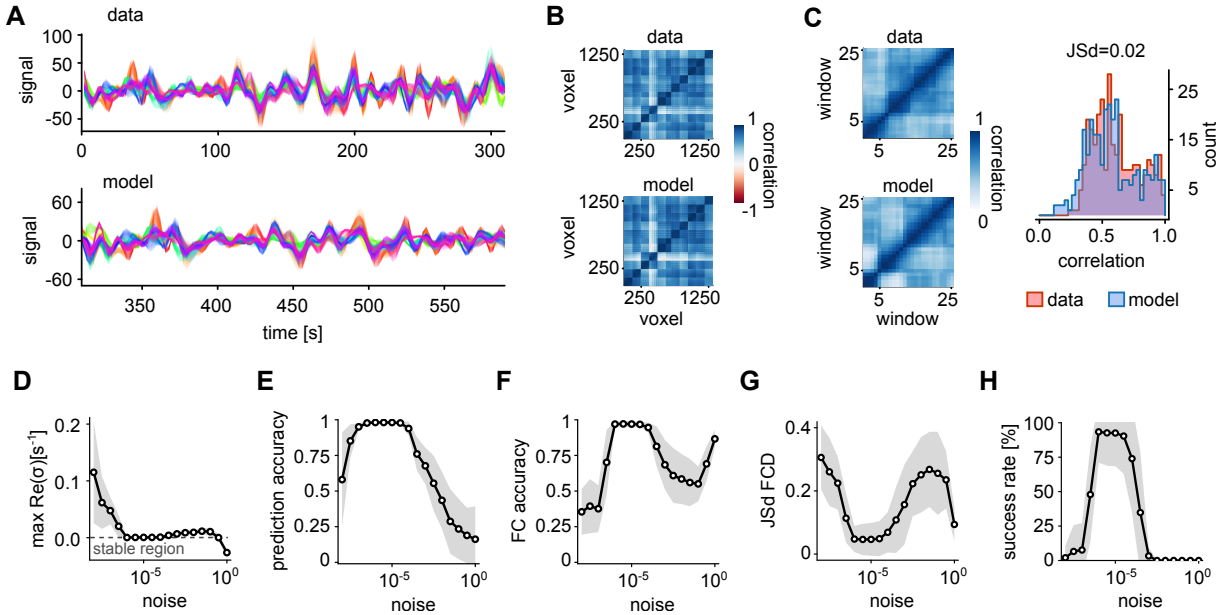

Figure 3: **Endogenous noise in IRNN is needed to replicate brain activity**. (**A**) Comparison of actual brain activity (top) with the corresponding output from a trained autonomous IRNN over the subsequent prediction interval (bottom). (**B** and **C**) Analysis of functional connectivity (FC) and functional connectivity dynamics (FCD) computed from the activity depicted in (A). Right, distribution of FCD values for IRNN-generated and exeprimental time series with the related Jensen-Shannon distance. (**D**, **E**, **F** and **G**) Features and performance statistics of inferred IRNNs across subjects varying the amplitude of endogenous noise: maximum pole real part (D), correlation between data and model prediction (E), FC correlation with actual data (F), and Jensen-Shannon distance between FCDs (G). (**H**) Success rate in reproducing above-threshold performances of the features in panels (E-G) (see main text for details).

in our digital twin, we incorporated an endogenous noise into each network unit. This modification introduces a generative mechanism that sustains activity over time. More specifically, the autonomous dynamics of the inferred (i.e., learned) IRNN dynamics is a multi-variate Ornstein-Uhlenbeck process

$$\tau d\boldsymbol{r} = (\tilde{\boldsymbol{W}} - \boldsymbol{I})\boldsymbol{r}dt + d\boldsymbol{\eta}\,, \tag{9}$$

where $\boldsymbol{\eta}(t)$ is the the vector of Gaussian white noise each unit of the network independently receives ($\langle d\eta_j(t) d\eta_k(t')\rangle = \gamma^2 \delta_{jk}\delta(t-t')$) with noise intensity γ. Note that, since the system is linear, this formulation does not change the theoretical framework derived above; in this case, it applies to the expectation values $\mathbb{E}[\boldsymbol{r}]$ and $\mathbb{E}[\boldsymbol{\omega}] = \boldsymbol{W}^{\text{out}}\mathbb{E}[\boldsymbol{r}]$.

To validate the effectiveness of this stochastic version of the IRNN in accurately replicating brain activity, we numerically integrated the network dynamics for an additional 387.5 seconds (Fig. 3A). We then compared the functional connectivity (FC) – defined as the correlation between BOLD signals of all possible voxel pairs – calculated from this simulation to that obtained from the experimental data (Fig. 3B). Furthermore, the temporal evolution of the FC, referred to as functional connectivity dynamics (FCD), closely aligned with the experimental observations. This similarity was quantified by the Jensen-Shannon distance (JSd) between the distribution values of the associated matrices (Fig. 3C).

Finally, we conducted a comprehensive analysis of the optimal intensity of noise across all subjects (Fig. 3), focusing on the accuracy of BOLD signal forecasting and the IRNN capability to reproduce both FC and FCD. To evaluate the quality of these aspects simultaneously, we introduced a success rate defined by applying fixed thresholds to the performance metrics: accuracy $> 80\%$, FC correlation $> 85\%$ and $\text{JSd} < 0.1$. Additionally, we assessed the stability of the inferred IRNN by examining the maximum real part of the poles σ_k (Fig. 3D). This analysis enabled us to identify a common optimal intensity γ of endogenous noise across subjects with a standard deviation of 10^{-6} (Fig. 3h). Interestingly, the existence of an optimal noise displays some similarities with the resonance phenomenon found in other machine-learning studies (39), which in our case robustly emerged in all the subjects.

Dynamical properties of the digital twin

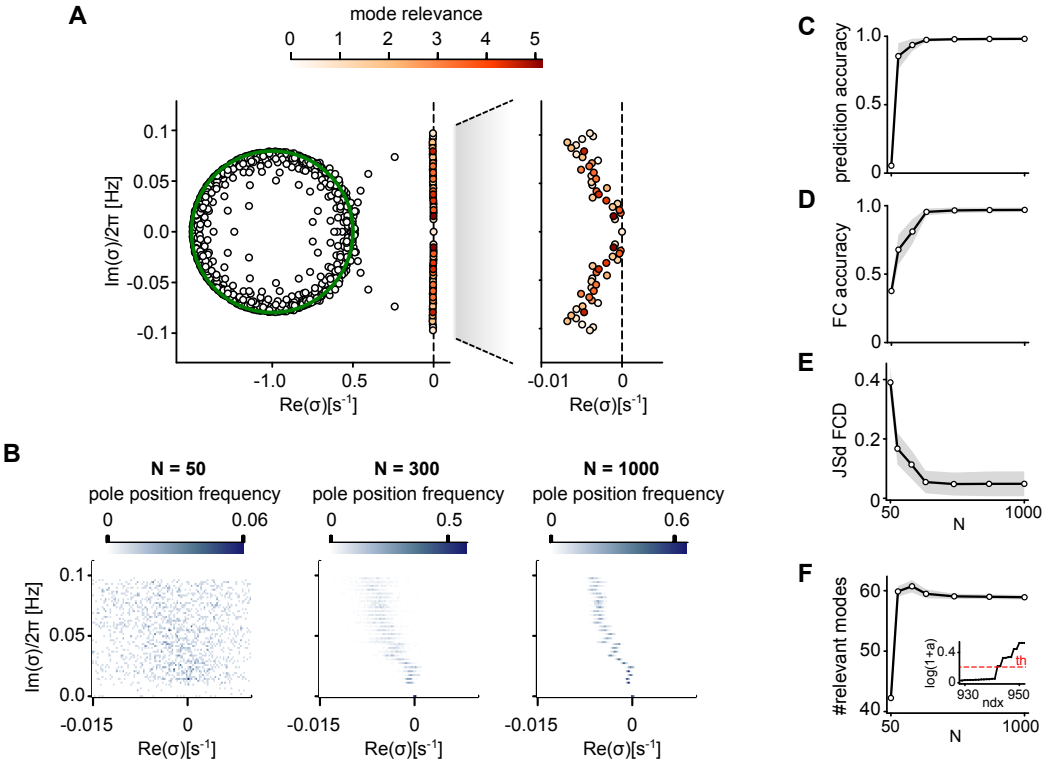

Figure 4: **Convergence of the mode decomposition with increasing IRNN size.** (**A**) Mode decomposition of a single cortical area for an example subject. Eigenvalues of $\tilde{\boldsymbol{W}}$ are color-coded as in Fig. 1E. (**B**) Distribution of eigenvalues for three distinct network sizes N, analyzed across 100 randomly sampled IRNNs. (**C**) Number of relevant modes (i.e., with relevance greater than 5%). Inset, sorted mode relevance and related threshold (dashed line). (**D**) Accuracy of forecasted activity. (**E** and **F**) Reproduction fidelity of FC (E) and of FCD (F). Gray shadings and black lines, standard deviation and mean across subjects, respectively.

We showed that RC with IRNNs can effectively reproduce resting state fMRI activity. Each simulated voxel can be described as a decomposition of linearly evolving modes. The dynamical properties of this decomposition can be effectively represented by the poles σ_k of the Laplace transform of the inferred IRNN, which are the eigenvalues $\tilde{\lambda}_k$ of the learned synaptic matrix $\tilde{\boldsymbol{W}}$. This spectrum of eigenvalues is shown in Fig. 4A for an example subject. For each of the modes, the relevance is averaged across all voxels and it is color coded, showing a pattern similar to what seen in Fig. 1E where only a subset of poles are moved towards the imaginary axis. These poles are the most relevant and display a specific organization like a rotated parabola. The modes with the highest relevance (i.e., the darkest) appear to be distributed at specific frequencies falling into the range of infra-slow oscillations (< 0.1 Hz), a typical footprint of the resting state and of the unconscious brain activity (40, 41, 42, 43). As a representative example, in Fig. 4A, the most relevant mode is one of the most persistent, characterized by a relatively small $\mathrm{Re}\,\sigma_k$, and exhibits a resonant frequency of approximately 0.02 Hz.

Subsequently, we investigated the stability of this representation exploring how the spectrum of $\tilde{\boldsymbol{W}}$ and the mode relevance change according to the number N of units in the IRNN. As shown in Fig. 4B, the density of eigenvalues in the complex plane $(\mathrm{Re}\,\sigma, \mathrm{Im}\,\sigma)$ becomes less and less scattered with increasing network size. Besides, all the performance measures converged to fixed values corresponding to a high quality of the data replicate by the digital twin (Fig. 4C-E), proving that a finite size IRNN can reach asymptotically high performances. Intriguingly, as convergence is approached the average number of relevant modes stands at about 60 independent components (Fig. 4F). This number is significantly higher than the number of experimental PCs provided as input to the IRNN (44 on average). Thus, brain activity replicated by the inferred IRNN appear to live in a latent state space whose dimensionality is larger than the one determined by its observation.

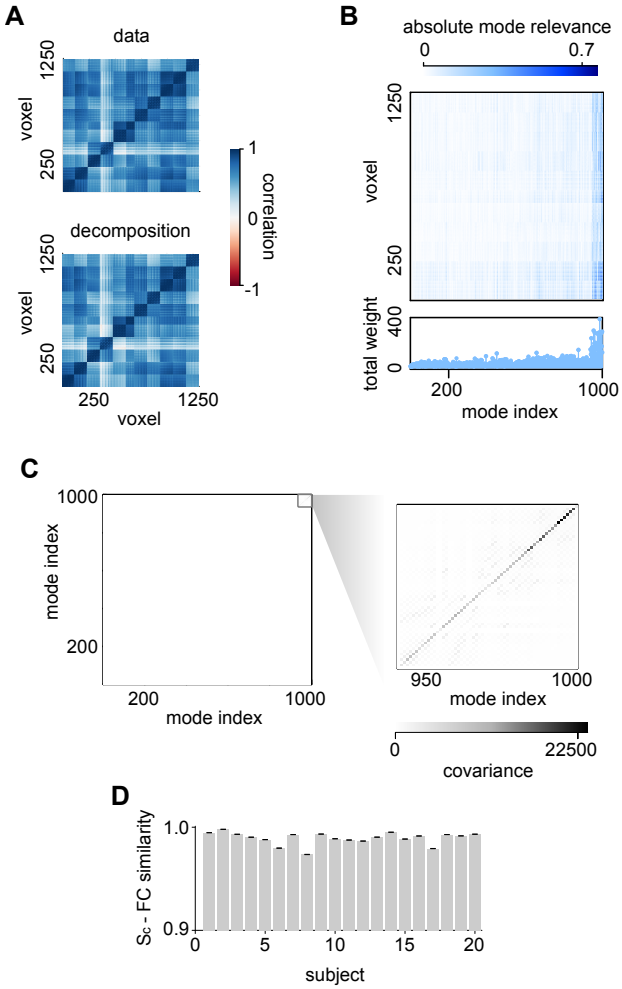

Figure 5: **Low-dimensional mode decomposition explains functional connectivity.** (**A**) Functional connectivity of an example subject (top) and cosine similarity (bottom) between rows of the matrix $\boldsymbol{Z}$ mapping eigenmode projections to voxel-wise BOLD activity from the inferred IRNN with $N = 1000$ units (see main text). (**B**) Absolute mode relevance $|Z_{jk}|$ for the same IRNN together with the total weight $\sum_j |Z_{jk}|$ of each of the N modes (bottom). Modes are sorted according to the real part of the associated eigenvalue (Re $\tilde{\lambda}_k$). (**C**) Covariance matrix of the eigenmode projections $v_k(t)$ (see Eq. (7)) during the extended autonomous phase (387.5 seconds post learning in Fig. 3A-bottom). Right, zoomed-in view on the slow and most relevant modes. (**D**) Correlation between the upper diagonal elements of the FC and the cosine similarity ($\mathrm{S_C}$) of the rows of $\boldsymbol{Z}$ for each subject. Averaged over 100 noise realizations. Error bars, SEM.

Functional connectivity illustrated in Fig. 5A-top for an example subject is then fully replicated by the inferred (and stochastic) IRNN. This effective copy occurs even though the dimensionality of the latent state space of the digital twin is significantly smaller than the number of voxels encompassing the examined language network. To gain a deeper understanding of this result, we 'opened the box' by examining the linear transformation $\boldsymbol{Z} \equiv \boldsymbol{P}^{-1}\boldsymbol{\Xi} \in \mathbb{R}^{L\times N}$. According to Eq. (7), this matrix facilitates the mapping of the N eigenmode projections $\boldsymbol{v}(t)$ onto the L-dimensional voxel-wise BOLD activity, represented as $\boldsymbol{x}(t) = \boldsymbol{Z}\boldsymbol{v}(t)$. The absolute values of the elements in this matrix are displayed in Fig. 5B for the same subject. Unsurprisingly, the most significant contributions arise from the slowest (rightmost) modes, specifically those with the largest Re $\tilde{\lambda}_k$. Additionally, by measuring the covariance between the eigenmode projections $v_k(t)$, it becomes evident from Fig. 5C that they are largely uncorrelated. Only about 60 of the slowest modes exhibit significant variability, as indicated by the dark diagonal elements. This empirical evidence suggests that functional connectivity can be estimated directly using the cosine

similarity $\mathrm{S_C}$ between the rows of the matrix $\boldsymbol{Z}$ (see Materials and Methods):

$$\mathrm{FC}_{jk} \approx \mathrm{S_C}\left(\langle Z_j|\,,\langle Z_k|\right)\,. \tag{10}$$

In Fig. 5A-bottom, the similarity matrix is presented, showcasing a remarkable overlap with the experimental functional connectivity $\mathbf{FC}$. It is important to note that in computing $\boldsymbol{Z}$, we considered only the columns of $\boldsymbol{\Xi}$ associated with the relevant modes. This further confirms that the remaining $N-60$ degrees of freedom are nearly irrelevant.

IRNN as a proxy to characterize subjects and brain areas

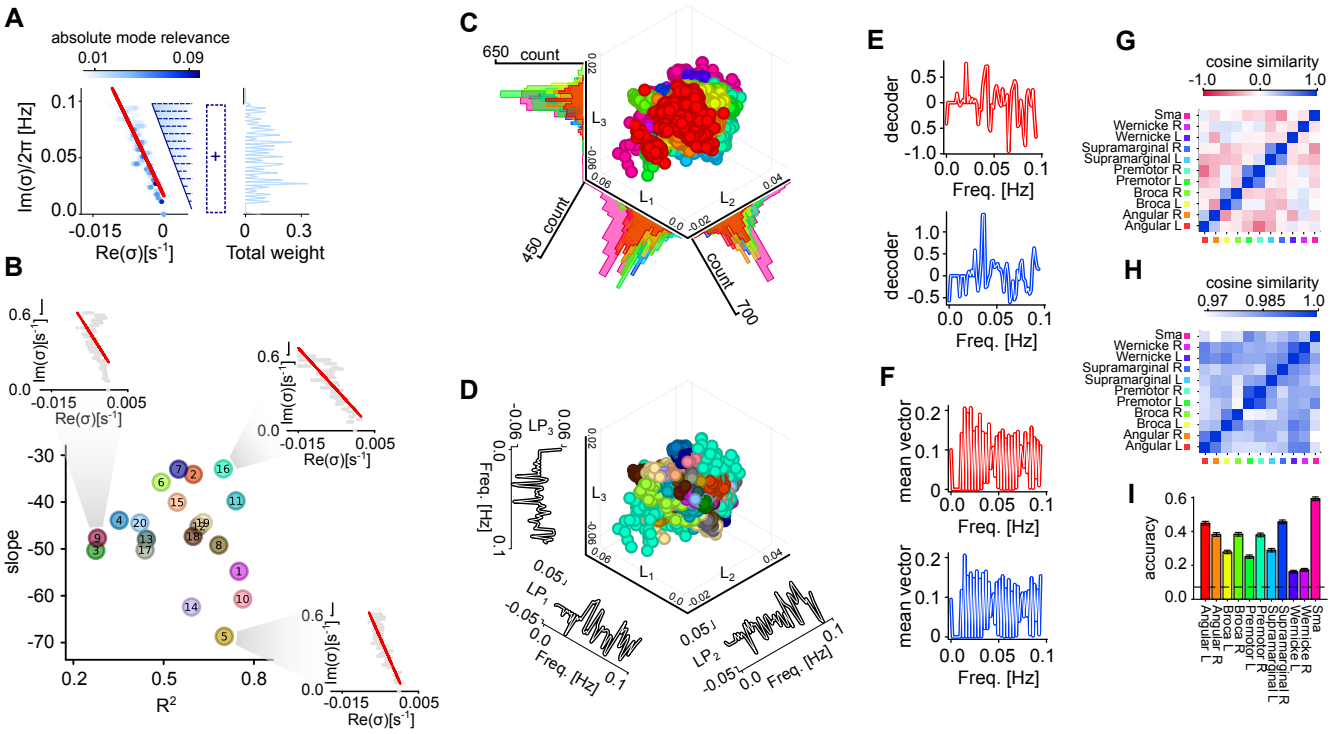

Figure 6: **Characterization of subjects and brain areas.** (**A**) Left, Spectrum of relevant eigenmodes for a sample subject. A linear regression (red line) is performed between the real and imaginary parts of the poles, σ_k, with the slope serving as the first feature that characterizes the subject. The coefficient of determination, R^2, from this regression represents the second feature. Right, Frequency vector, i.e., the histogram of $\mathrm{Im}\,\sigma_k$ each counted with multiplicity given by the averaged absolute mode relevance across voxels in a single brain area: $\langle|\Xi_{jk}|\rangle_{j\in\mathrm{BA}}$. (**B**) Slope of linear regression and R^2 of the 20 subjects in the dataset. Insets: eigenmode spectra of three representative subjects with the related linear regressions (red lines, see main text for details). (**C**) First three latent variables derived from the linear discriminant analysis (LDA) applied to the frequency vectors of 11 brain areas across all 20 subjects. The densities of circles for each latent variable are also plotted. For each subject, 100 IRNNs are inferred by sampling different endogenous noise in their stochastic dynamics, resulting in a total of 22,000 circles. Colors code brain areas according to Fig. 2A. (**D**) Same as (C) but with circles colored according with the subject color as in (B). Composition of the three latent variables LP_1, LP_2, LP_3 are show on the side of their respective axes. (**E** and **F**) Linear decoders (E) discriminating the different areas in the frequency-vector space and average frequency vectors (F) for two sample brain areas (red and blue). (**G** and **H**) Cosine similarity between linear decoders (G) and between frequency vectors (H) per area averaged across subjects. (**I**) Accuracy of decoders along the training set. Dashed black line highlights the random chance level separation.

Given that inferred IRNNs reliably replicate observed brain activity and their modal decomposition defines their dynamical properties, a critical question emerges: Can IRNNs serve as proxies for understanding the similarities

and differences among subjects and cortical areas? To address this, we characterized the spectrum of relevant eigenvalues by focusing on two key features: the linear relationship between their real and imaginary parts, and the overall significance of each oscillation frequency, as illustrated in Fig. 6A. The slope of this linear relationship indicates the correlation between decay times and oscillation frequencies, while the coefficient of determination (R^2) assesses the goodness of fit for this linear model. In this two-dimensional parameter space, subjects are systematically distributed among three distinct regions (Fig. 6B). Those with high R^2 values display either steep or shallow slopes, whereas subjects with low R^2 values indicate a poor linear fit and suggest the presence of persistent isolated high-frequency oscillatory modes. Within this representational space, spectra characterized by steeper slopes (located at the bottom) are associated with poles that are closer to the imaginary axis. This suggests that those lRNNs (and consequently subjects) exhibit longer relaxation time scales, which may be linked to brain activity approaching a critical point where the resting state could become unstable and displaying a more complex dynamics.

Shifting our focus to the frequency vectors derived from the inferred lRNN, we sought statistical regularities that reveal invariant features associated with individual brain areas and subjects. The frequency vector for a given brain area in a subject is represented by the histogram of Im σ_k from the corresponding lRNN, each counted with multiplicity proportional to its absolute mode relevance averaged across the voxels in that area: $\langle|\Xi_{jk}|\rangle_{j\in\mathrm{BA}}$. We then conducted a linear discriminant analysis (LDA) in the frequency-vector space to identify patterns that effectively discriminate between different brain areas (Fig. 6C). Despite originating from distinct subjects, the resulting data points clustered according to their respective brain areas. It is important to note that the compositions of the latent variables (L_1, L_2, and L_3) correspond to specific frequency patterns (Fig. 6D). Among these, the pattern associated with L_1 appears most relevant, as data points on this linear manifold are distinctly clustered based on the brain area of belonging (Fig. 6C).

To further isolate the frequencies at which each brain area exhibits unique dynamics, we computed the corresponding linear decoders (see Materials and Methods). The elements of the coefficient vectors characterizing each decoder are assigned high absolute values if the associated frequency plays a crucial role in distinguishing that area from others (Fig. 6E). The resulting patterns are markedly distinct from the average frequency vector, which, in contrast, reveals a continuum of relevant frequencies (Fig. 6F). In fact, the frequency vectors show significant similarities, making it challenging to differentiate between areas even if they are located in different hemispheres (Fig. 6H). In contrast, the decoders exhibit a lower degree of similarity, except when comparing the corresponding areas of both hemispheres (Fig. 6G). Interestingly, there is an exception to this trend: the decoder for the right Wernicke area appears to highlight frequencies that differ from those identified for the same area in the left hemisphere. Further insights into brain lateralization are revealed by examining the accuracy of the decoders. Specifically, it is more difficult to linearly recognize left hemisphere patterns by this method for the Broca, Premotor, and Supramarginal areas (Fig. 6I). It is worth noting that Wernicke is the most difficult to recognize.

In short, these findings demonstrate that inferred lRNNs can serve as powerful proxies to characterize the unique dynamics of subjects and areas of the brain, capturing their specific dynamical footprints.

Discussion

The successful application of linear recurrent neural networks within the reservoir computing framework represents a significant methodological contribution to model resting-state fMRI data. Our results show that lRNNs effectively capture brain dynamics, suggesting that resting-state activity can be approximated as a linear system in a high-dimensional space, supporting recent studies advocating for linear models in this context (44). Our approach provides valuable and easily accessible insights into the functional connectivity structure as similarity patterns between vectors associated with the eigenmode decomposition of the inferred lRNN. This low-dimensional description of the whole brain activity moves the focus on the modes subspace offering a normative framework where brain areas and inter-subject comparison is more straightforward.

Notably, the mode dynamics described in Eq. (6) provides an equivalent linear representation for the observables of the system under investigation – namely, the single-subject voxel-wise BOLD activity. The resulting vector of all time derivatives (see Materials and Methods) serves as an alternative descriptor of the system's state (45), potentially encapsulating the trajectory evolution at any given time. This concept is reminiscent of the dynamic mode decomposition (DMD) with Hankel matrix (46, 47), but with a significantly reduced dimensionality constraint. Unlike the DMD approach, which requires estimating a number of parameters that scales quadratically with the

number M of observables – specifically $(M \times T)^2$ parameters (where T denotes the number of delayed time samples) – our reservoir computing approach only necessitates inferring order N readout weights, corresponding to the size of the IRNN. This result also stems from considering an optimal amount of memory-less fluctuations integrated into the IRNN units' endogenous activity. This feature enables effective modeling of non-informative fast fluctuations in BOLD activity without requiring inference of additional high-frequency, fast-decaying modes, thereby reducing the dimensionality of the latent state space visited by the deterministic components of the inferred IRNNs. This characteristic makes the presented method particularly advantageous and flexible, allowing the inference of IRNN-based digital twins at the single-subject level using a limited amount of experimental data.

The reported results have the potential to open up several avenues for future research. Indeed the possibility offered to represent the single-subject as a point into a low-dimensional map – i.e., a landscape (Fig. 6B) – allows in principle to follow in time its drift in longitudinal studies like those focused on neurodegenerative diseases (48) or aging (49, 50, 51). Here, the novelty is the fully data-driven nature of the approach to build such landscape of subjects compared to those relying on model-driven inferences (52). Characterizing the trajectory followed in this landscape by each subject could provide valuable insights into the alterations in functional connectivity and dynamics. For instance, the emergence of unstable modes may be linked to transitions into pathological states of a specific neurological disease. By perturbing the system with specific and personalized rehabilitative or pharmacological treatments that aim to suppress these "pathology-related" modes, we may gain a means to intervene and reduce the occurrence of such transitions, ultimately helping to tailor optimal therapeutic strategies (28, 53).

In line with this, as in classic digital twins, the inferred IRNN can incorporate any additional available information about the system simply as new input. An external stimulation can potentially act as a selector for different dynamical regimes, allowing for two distinct representations with a single network (54). The functional connectivity itself could be a proxy for state transition. Indeed, a significant change in functional connectivity might indicate that the decomposition has also changed, suggesting that the observed system is close to a different equilibrium point (55). These changes can be tracked following the distribution of poles that characterize the autonomous dynamics of the inferred linear RNN. This is the case for transitions in global brain states – such as those governing the sleep-wake cycle – that arise from network destabilization, a hallmark of criticality (56, 57, 58). Under these conditions, pole distributions correlate with longer relaxation time scales and steeper slopes (as illustrated in Fig. 6B), possibly underlying previously observed state-dependent spectral signatures (59). This metastable dynamics results in broader excursions within the latent state space of neural activity, making predictions of future BOLD activity more challenging for the inferred IRNN. The quality of IRNN predictions can then serve as a proxy for dynamical stability, akin to findings in intracortical local field potentials (LFPs) of macaque monkeys during anesthesia-induced transitions between wakefulness and unconsciousness (60).

While this study demonstrates the efficacy of IRNNs in modeling resting-state brain activity, several limitations must be acknowledged. The assumption of linearity, although supported by our results, may not fully capture the complexity of brain dynamics under all conditions. Nonlinear phenomena, transient states, and the influence of external stimuli could require more sophisticated modeling approaches (61, 14). Alternatively, the apparent suitability of resting-state activity being well represented by a stationary multivariate Ornstein-Uhlenbeck process (i.e., Brownian motion) (62, 63) may reflect an intrinsic limitation of the BOLD signal in conveying detailed neural dynamics. Indeed, this hemodynamic-related signal acts as a lumped observable that inevitably linearizes the spiking activity of relatively large neuronal assemblies (64, 44). The coexistence of endogenous noise with the relatively high-dimensionality of the BOLD time series can give rise to a multivariate Brownian motion exhibiting a rich repertoire of restless spatiotemporal modes (63). These are captured in our IRNN by relatively wide distributions of poles. Within this linear stochastic system, noise spontaneously excites its normal modes, eliciting co-fluctuating patterns closely linked to the functional connectivity measured in rs-fMRI. Interestingly, this structured stochasticity may, at least in principle, exhibit similarities to turbulence (65, 66).

However, the existence of a single fixed-point around which all voxel activities fluctuate, imposes a rather strong constraint on the brain's ability to perform information processing. Indeed, this computational capability is thought to arise from the complex unfolding of neural trajectories across a landscape rich in saddles and metastable states (67, 68). Yet, IRNNs cannot represent the latent state space of a multistable system. This raises an important question: how can we reconcile the evidence that resting-state fMRI activity is well described by stochastic IRNNs with the necessity of nonlinear dynamics for brain computation? A tentative answer is that, during rest, endogenous fluctuations have a magnitude comparable to the deterministic excursions of neural activity associated with relaxation dynamics. Consequently, the nonlinear components of brain activity may only

become apparent when the system is pushed far from equilibrium – such as when the brain engages in cognitive functions like motor or perceptual decision-making. A promising direction for future work would then be to test this hypothesis by investigating to what extent an effective lRNN can be inferred from BOLD time series recorded while subjects engage in such cognitive tasks.

Materials and Methods

Dataset description

Subjects. The study included 20 healthy, right-handed subjects (mean age $\pm$ SD = 37 ± 12; 7 females, 13 males) with no history of neurological disorders. The study was approved by the Institutional Review Board at Memorial Sloan Kettering Cancer Center, and informed consent was obtained from each participant. During the resting-state condition, subjects were instructed to lie in the scanner, keep their eyes open, try to think of nothing in particular, and maintain fixation on a central cross on the screen.

MRI methods. A GE 3T scanner (General Electric, Milwaukee, Wisconsin, USA) and a standard quadrature head coil was employed to acquire the MR images. Functional images covering the whole brain were acquired using a (T2*)-weighted imaging sequence sensitive to blood oxygen level-dependent (BOLD) signal (repetition time, TR/TE = 2500/40 ms; slice thickness = 4.5 mm; matrix = 128 $\times$ 128; FOV = 240 mm; volumes = 160). Functional matching axial T1-weighted images (TR/TE = 600/8 ms; slice thickness = 1 mm) were acquired for anatomical co-registration purposes.

Data preprocessing. Functional MRI data were processed and analyzed using the software program Analysis of Functional NeuroImages (AFNI; Cox, 1996). Head motion correction was performed using 3D rigid-body registration. The first volume was selected to register all other volumes. The first volume was chosen because it was acquired before the anatomical scan. Both task and resting state fMRI scans were monitored using a real time post-processing software BrainWave (BrainWave RT, Medical Numerics, Germantown, MD) to monitor brain activity and the head motion. For subjects showing severe head motion over time, generally, the scan is repeated. For small head motion (less than 2 voxel size), a motion correction algorism (iterated linearized weighted least squares) considering 3 translation and 3 rotation parameters against a reference volume was applied. The obtained 6 parameter time courses were also integrated in the statistical analysis to regress out residual motion-correlated artifactual voxels. Spatial smoothing was applied to improve the signal-to-noise ratio using a Gaussian filter with a 6 mm full width of half maximum. Corrections for linear trend and high-frequency noise were also applied. Resting-state data requested some more preprocessing steps. They were corrected for head motion by regressing head motion data and the first five principal components of the white matter and CSF signals. They were also detrended, demeaned, and band-pass filtered (frequency range 0.01-0.1 Hz). All fMRI data were registered to the standard space (Montreal Neurological Institute MNI152 standard map).

fMRI dimensionality analysis

Principal Component Analysis (PCA) is a widely used dimensionality reduction technique used to transform a dataset into a lower-dimensional space while preserving most of the information in the original data (69). In the context of the given problem, PCA is applied to the set of BOLD signals to reduce redundancy and keep only the relevant information.

The PCA-based dimensional reduction of BOLD is carried out based on the initial 280 s (learning period) of each subject's data. Referring to Fig. 2, the number of principal components taken into account are those capturing 99.5% of the variance in the time series of each area. The resulting transformation is represented by a block matrix $\boldsymbol{P}_0$, where each block's columns $\boldsymbol{p}_i$ correspond to the eigenvectors of the covariance matrix of the respective cortical area:

$$\boldsymbol{P}_0 = \begin{bmatrix} \boldsymbol{p}_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \boldsymbol{p}_{11} \end{bmatrix} \tag{11}$$

The output undergoes orthogonalization through an additional PCA step ($\boldsymbol{P}_1$) utilizing all components. Subsequently, normalization is performed by applying the diagonal matrix $\boldsymbol{N}$, which contains the reciprocals of the maximum values of the compressed signals observed during the learning period. This approach ensures that all currents enter the reservoir with the same maximum strength. The final transformation matrix $\boldsymbol{P}$ is computed as:

$$\boldsymbol{P} = \boldsymbol{N}\boldsymbol{P}_1\boldsymbol{P}_0 \tag{12}$$

While the antitransformation matrix $\boldsymbol{P}^{-1}$ is calculated as:

$$\boldsymbol{P}^{-1} = \boldsymbol{P}_0^T\boldsymbol{P}_1^T\boldsymbol{N}^{-1} \tag{13}$$

The data is effectively reduced and then restored, ensuring the preservation of essential information while reducing dimensionality for the analysis. The decomposition for voxel activity is derived by applying $\boldsymbol{P}^{-1}$ to the decomposition matrix for the actual network inputs.

Koopman Operator

The discrete-time Koopman operator $\mathcal{K}^1 : \mathcal{F} \to \mathcal{F}$ is a linear operator defined in an infinite-dimensional space $\mathcal{F}$ of observables of the system state (70, 19, 20). The operation of the Koopman operator on an observable g is described by the equation:

$$\mathcal{K}^1 g(\boldsymbol{x}_t) = g(\boldsymbol{x}_{t+1})\,. \tag{14}$$

Here, $\boldsymbol{x}$ represents a point in the state space. The eigenfunctions of the Koopman operator are potential observables of the system themselves, and have the peculiar property of evolving linearly in time:

$$\mathcal{K}^1\psi_k(\boldsymbol{x}_t) = \psi_k(\boldsymbol{x}_{t+1}) = \lambda_k^1\psi_k(\boldsymbol{x}_t)\,. \tag{15}$$

These eigenfunctions $\psi_k(\boldsymbol{x}_t)$ constitute a basis of the space $\mathcal{F}$, enabling the representation of all the possible observables of the system as a linear combination of linearly evolving functions.

In general there exists a family of Koopman operators, $\mathcal{K}^t$, that advances a function forward by a time t:

$$\mathcal{K}^t g(\boldsymbol{x}) = g(\mathcal{S}^t(\boldsymbol{x}))\,. \tag{16}$$

Here, $\mathcal{S}^t$ denotes the time flow operator, with $\boldsymbol{x}(t) = \mathcal{S}^t[\boldsymbol{x}(0)]$. The infinitesimal generator of the Koopman operator family $\mathcal{L}g = \lim_{t\to 0}\frac{\mathcal{K}^t g - g}{t}$ is the Lie operator which evaluates the temporal change of observables:

$$\mathcal{L}g(\boldsymbol{x}(t)) = \frac{d}{dt}g(\boldsymbol{x}(t)) \tag{17}$$

The family of Koopman operators can then be expressed in this term $\mathcal{K}^t = e^{\mathcal{L}t}$.

Reservoir computing

The reservoir computing approach involves a RNN with fixed and random internal couplings (namely the 'reservoir') whose state is fed forward to a second set of 'readout' units (31, 4). In this framework the required 'echo-state property' spontaneously emerges from the RNN collective dynamics. The units composing the RNN are intended to model homogeneous and local neuronal assemblies of a cortical network (71, 72, 73). Each of the N units in the network has activity state $r_j(t)$ evolving in time as

$$\tau\dot{r_j} = \Phi(h_j) - r_j \tag{18}$$

with $j \in [1, N] \subset \mathbb{Z}$. The decay time constant τ and the sigmoidal activation function $\Phi(h_j)$ is the same for all units. Here we consider the limiting case of weak recurrent coupling where the activation function can be approximated to a linear function: $\Phi(h_j) = h_j$. The synaptic input h_j is the weighted sum

$$h_j(t) = \sum_{k=1}^{N} W_{jk}r_k(t) + h_j^{\text{in}}(t)$$

where W_{jk} are the elements of the synaptic matrix (i.e., the internal couplings) $\boldsymbol{W} \in \mathbb{R}^{N\times N}$, and $h_j^{\rm in}(t)$ is the external input received by the unit j. In reservoir computing, the external input is driven by be the measured observable ω of the inspected system, which is fed into the network using random synaptic couplings: $h_j^{\rm in}(t) = W_j^{\rm in}\,\omega(t)$. Due to the high dimensionality of the trajectories in the RNN state space, a simple linear transformation represented by $\boldsymbol{W}^{\rm out}$ is often sufficient to accurately map the RNN state into a desired output time series. More precisely, given $\boldsymbol{R} \in \mathbb{R}^{N\times T}$ whose k-th row represent the time series of the k-th unit during a learning period lasting T time steps, and $\boldsymbol{\Omega} \in \mathbb{R}^{M\times T}$ has rows given by the M observables $\omega_k(t)$ to be replicated, the linear map is computed by a ridge regression:

$$\boldsymbol{W}^{\rm out} = \boldsymbol{\Omega}\boldsymbol{R}^T(\boldsymbol{R}\boldsymbol{R}^T + \beta\boldsymbol{I})^{-1} \tag{19}$$

where β is a regularization parameter and $\boldsymbol{I} \in \mathbb{R}^{N\times N}$ the identity matrix. The target output coincides with the external input. Then, the resulting linear map can be used to simulate the external stimulation and predict the future steps of the data. This is equivalent to update the synaptic matrix of the reservoir. In particular $\boldsymbol{W}^{\rm in}\boldsymbol{\omega}(t) \approx \boldsymbol{W}^{\rm in}\boldsymbol{W}^{\rm out}\boldsymbol{r}(t)$ leading to the autonomous system:

$$\tau\dot{\boldsymbol{r}} = \boldsymbol{\Phi}\left[\tilde{\boldsymbol{W}}\boldsymbol{r}\right] - \boldsymbol{r}\,, \tag{20}$$

where $\tilde{\boldsymbol{W}} = \boldsymbol{W} + \boldsymbol{W}^{\rm in}\boldsymbol{W}^{\rm out}$ is the updated synaptic matrix.

Linear recurrent neural networks (IRNN) and autoregressive models

In the established framework of reservoir computing with linear RNNs (IRNN, i.e., with a linear activation function) by setting as initial condition $r_k(-\infty) = 0$, the system Eq. (18) has the following solution:

$$\boldsymbol{r}(t) = \sum_{k=1}^{M}\int_{-\infty}^{t}\omega_k(t')e^{(\boldsymbol{W}-\boldsymbol{I})(t-t')/\tau}\,|W_k^{\rm in}\rangle\,dt' \tag{21}$$

As above $\omega_k(t)$ denotes the k-observable measured from the inspected system. Assuming a time t^* exists after which a matrix $\boldsymbol{W}^{\rm out}$ maps the network state to the input data $\boldsymbol{\omega}(t) = \boldsymbol{W}^{\rm out}\boldsymbol{r}(t)$, we can rewrite the above equation as:

$$\boldsymbol{\omega}(t) = \sum_{k=1}^{M}\int_{-\infty}^{t}\omega_k(t')\boldsymbol{W}^{\rm out}e^{(\boldsymbol{W}-\boldsymbol{I})(t-t')/\tau}\,|W_k^{\rm in}\rangle\,dt' \tag{22}$$

By substituting $s = t - t'$, we obtain an autoregressive description of the signals:

$$\omega_j(t) = \sum_{k=1}^{M}\int_{0}^{\infty}G_{jk}(s)\omega_k(t-s)ds \tag{23}$$

Here, $G_{jk}(s) = \boldsymbol{W}_j^{\rm out}e^{(\boldsymbol{W}-\boldsymbol{I})s/\tau}\boldsymbol{W}_k^{\rm in}$. These kernels can be expressed in terms of the synaptic matrix spectrum. The diagonal representation of the exponential matrix is given by:

$$e^{(\boldsymbol{W}-\boldsymbol{I})s/\tau} = \sum_{n=1}^{N}e^{(\lambda_n-1)/\tau}\,|Q_n\rangle\,\langle Q_n^{-1}|\ . \tag{24}$$

In this representation, λ_n denotes the eigenvalues and $|Q_n\rangle\langle Q_n^{-1}|$ the outer product between the right and left eigenvectors of $\boldsymbol{W}$. This representation allows us to rewrite the kernel as the weighted sum over exponentially decaying components:

$$G_{jk}(s) = \sum_{n=1}^{N}J_{jk}^{(n)}e^{(\lambda_n-1)s/\tau}\,, \tag{25}$$

where the tensor of weights is defined by the set of rank-1 matrices $\boldsymbol{J}_n = \boldsymbol{W}^{\rm out}\,|Q_n\rangle\,\langle Q_n^{-1}|\,\boldsymbol{W}^{\rm in}$.

IRNN and the Laplace transform of system state

Assuming the existence of a time t^* such that a matrix $\boldsymbol{W}^{\text{out}}$ maps the IRNN state into the system observable given as input $\omega(t) = \boldsymbol{W}^{\text{out}}\boldsymbol{r}(t)$, the open- and closed-loop dynamics are in principle equivalent (31, 30). This property allows us to link the kernel of the autoregressive model (closed-loop dynamics) to the spectrum of the learned synaptic matrix $\tilde{\boldsymbol{W}}$, and consequently, to the resulting finite Koopman approximation, as we will see in the following.

Open-loop Representation

The autoregressive formula expressed in Eq. (3) can be reformulated as:

$$\omega_j(t) = \sum_{k=1}^{M} \int_0^t G_{jk}(t-t')\omega_k(t')dt' + \omega_j(0)\,. \tag{26}$$

The convolution theorem for the Laplace transform leads to:

$$\text{LT}\{\omega_j\}(\sigma) = \sum_{k=1}^{M} \text{LT}\{G_{jk}\}(\sigma)\text{LT}\{\omega_k\}(\sigma) + \frac{\omega_j(0)}{\sigma} \tag{27}$$

Here, $\text{LT}\{G_{jk}\} = \int_0^\infty G_{jk}(t)e^{-\sigma t}dt$ denotes the Laplace transform of the kernel that is

$$\text{LT}\{G_{jk}\}(\sigma) = \sum_{n=1}^{N} \frac{J_{jk}^{(n)}}{\sigma - \frac{\lambda_n - 1}{\tau}}\,. \tag{28}$$

Thus, in matrix formalism the Laplace transforms of the system observables result to be

$$\text{LT}\{\boldsymbol{\omega}\}(\sigma) = \frac{\boldsymbol{L}^{-1}(\sigma)}{\sigma}\boldsymbol{\omega}(0)\,, \tag{29}$$

where the matrix $\boldsymbol{L}(\sigma) \in \mathbb{R}^{M\times M}$ has elements $L_{ij}(\sigma) = \delta_{ij} - \text{LT}\{G_{ij}\}(\sigma)$.

Closed-loop representation

The closed-loop network evolves according to Eq. (5). Assuming a diagonalizable synaptic matrix $\tilde{\boldsymbol{W}}$, it can be factorized as $\tilde{\boldsymbol{W}} = \tilde{\boldsymbol{Q}}\tilde{\boldsymbol{\Lambda}}\tilde{\boldsymbol{Q}}^{-1}$, where $\tilde{\boldsymbol{Q}}$ is a matrix whose columns are the right eigenvectors of $\tilde{\boldsymbol{W}}$ and $\tilde{\boldsymbol{\Lambda}}$ is a diagonal matrix containing the associated eigenvalues. The spectral decomposition of the learned synaptic matrix leads to a decoupled set of linearly evolving projections $\boldsymbol{v}(t) = \tilde{\boldsymbol{Q}}^{-1}\boldsymbol{r}(t)$ whose N elements evolves independently as

$$v_n(t) = v_n(0)e^{(\tilde{\lambda}_n - 1)t/\tau}\,. \tag{30}$$

These variables serve as a basis to decompose the system observables $\omega_k(t)$ as the sum of linearly evolving modes, resembling the Koopman decomposition. Indeed, given $\boldsymbol{r} = \tilde{\boldsymbol{Q}}\boldsymbol{v}$ we can write

$$\omega_j(t) = \sum_{n=1}^{N} \Xi_{jn}v_n(t) = \sum_{n=1}^{N} \Xi_{jn}v_n(0)e^{(\tilde{\lambda}_n - 1)t/\tau}\,, \tag{31}$$

with matrix $\boldsymbol{\Xi} \in \mathbb{R}^{M\times N}$ defined as $\boldsymbol{\Xi} = \boldsymbol{W}^{\text{out}}\boldsymbol{Q}$ such that $\boldsymbol{\omega}(t) = \boldsymbol{\Xi}\boldsymbol{v}(t)$. The Laplace transform of the replicated observables is then

$$\text{LT}\{\omega_j\}(\sigma) = \sum_{n=1}^{N} \frac{\Xi_{jn}v_n(0)}{\sigma - \frac{(\tilde{\lambda}_n - 1)}{\tau}} \tag{32}$$

whose poles coincide with the eigenvalues of the estimated Koopman matrix.

Given the equivalence between the open and closed representations, $L_{ij}(\sigma)/\sigma$ shares the same poles, providing a link between the autoregressive modeling and the linear dynamic representation.

IRNNs and the spectrum of the Koopman operator

The closed-loop representation provides a decomposition of the signal as linear evolving modes. We can explicitly express the functional form of the basis starting from open loop dynamics. By applying $\tilde{\boldsymbol{Q}}^{-1}$ to both hand sides of Eq. (2) yields:

$$\boldsymbol{v}(t) = \sum_{k=1}^{M} \int_0^\infty \omega_k(t-s) \tilde{\boldsymbol{Q}}^{-1} e^{(\boldsymbol{W}-\boldsymbol{I})s/\tau} \left|W_k^{\text{in}}\right\rangle ds\,. \tag{33}$$

By resorting to the spectral decomposition in Eq. (24), the above eigenmode dynamics reduces to:

$$v_j(t) = \sum_{k=1}^{M} \sum_{n=1}^{N} \left[\tilde{\boldsymbol{Q}}^{-1} \left|Q_n\right\rangle \left\langle Q_n^{-1}\right| \boldsymbol{W}^{\text{in}}\right]_{jk} \int_0^\infty \omega_k(t-s) e^{(\lambda_n - 1)s/\tau} ds\,. \tag{34}$$

We can then expand in Taylor series the observables: $\omega_k(t-s) = \sum_{r=0}^{\infty} \frac{d^r \omega_k(t)}{dt^r} (-1)^r s^r / r!$. This allows to solve the above integral leading to a functional description of the estimated eigenmodes

$$v_j(t) = \psi_j\left[\boldsymbol{\omega}(t), \dot{\boldsymbol{\omega}}(t), ..\right] = \sum_{n=1}^{N} \sum_{k=1}^{M} A_{jk}^{(n)} \sum_{r=0}^{\infty} \frac{d^r \omega_k(t)}{dt^r} \left(\frac{\tau}{\lambda_n - 1}\right)^r\,. \tag{35}$$

The vector of all the derivatives $d^r\omega_k(t)/dt^r$ fully determine the state of the observed system (33) and ψ_j is the eigenfunction of the approximated Koopman operator depending on the full state of the system. The corresponding eigenvalue $\sigma_j = (\tilde{\lambda}_j - 1)/\tau$, and $A_{jk}^{(n)}$ is a tensor of complex weights defined by the rank-1 matrices $\boldsymbol{A}_n = \frac{\tau}{1-\lambda_n} \tilde{\boldsymbol{Q}}^{-1} \left|Q_n\right\rangle \left\langle Q_n^{-1}\right| \boldsymbol{W}^{\text{in}}$.

Stochastic IRNN

The deterministic neural network described can be extended to the stochastic case, where a white noise input stimulates the reservoir, maintaining it out of equilibrium as a continuous source of new energy. Each unit receives a total input defined by the equation:

$$h_j(t) = \sum_{k=1}^{N} W_{jk} r_k(t) + h_j^{\text{in}}(t) + \eta_j(t) \tag{36}$$

where $\eta_j(t)$ represents an Ornstein-Uhlenbeck process (colored noise) with a relatively small correlation time such that for the purpose of this work it can be considered as memory-less: $\langle \eta_i(t)\eta_j(t')\rangle = \gamma^2 \delta_{ij} \delta(t-t')$. An optimal standard deviation γ can be identified to maximize the accuracy of the linear map $\boldsymbol{W}^{\text{out}}$.

In the closed-loop formulation, the input can be expressed as:

$$h_j(t) = \sum_{k=1}^{N} \tilde{W}_{jk} r_k(t) + \eta_j(t) \tag{37}$$

Consequently, the associated Langevin equation for the reservoir is given by:

$$\tau d\boldsymbol{r} = -\boldsymbol{r} dt + \tilde{\boldsymbol{W}} \boldsymbol{r} dt + d\boldsymbol{\eta}\,. \tag{38}$$

In this context, the introduced autoregressive description and linear decomposition remain valid when considering the expected values of the system.

Functional Connectivity and Functional Connectivity Dynamics

Functional Connectivity (FC) is a measure of the degree of co-activation over time of different brain regions. Although FC does not allow to infer the directional flow of information among the nodes of the brain network,

it has proven to provide valuable insights into the functional organization of the central nervous system (74). Mathematically, the FC between two brain regions labeled A and B, is usually measured as the correlation function

$$\mathrm{FC}_{AB}(t;T) = \mathrm{cor}[X_A(t;T), X_B(t;T)]\,, \tag{39}$$

where cor is the Pearson correlation coefficient, and $X_A(t;T)$ and $X_B(t;T)$ are the BOLD signals from region A and B, respectively, in the time window $[t, t+T]$.

Temporal changes of functional connectivity are evaluated by computing the Functional Connectivity Dynamics (FCD) (75, 76). FCD is defined as the correlation between vectorized FC matrices evaluated in time-shifted chunks of BOLD activity:

$$\mathrm{FCD}_{A,B}(\delta t, t; T) = \mathrm{ucor}[\mathrm{FC}_A(t;T), \mathrm{FC}_B(t+\delta t;T)]\,, \tag{40}$$

where δt is a constant time shift, and ucor is the correlation between the vectors composed of the upper-diagonal elements of the two FC matrices involved.

Throughout the whole paper FC is computed with $t = 0$ s and $T = 387.5$ s, while FCD uses $T = 75$ s time shift δt multiples of 12.5 seconds.

Functional Connectivity in IRNN

Once a IRNN capable to reproduce the systems observables $\omega_j(t)$ is available, the elements of the FC matrix can be inferred directly from the IRNN parameters. The computed Functional Connectivity of the simulated signals depends on the covariance matrix estimated from time t over a time window of length T. Considering large time windows T and assuming a zero time average $\langle \omega_j(t) \rangle_t \simeq 0$, it holds:

$$\mathrm{cov}_{ij}(t;T) = \int_t^{t+T} \omega_i^*(t')\omega_j(t')dt' = \sum_{r,s} \Xi_{ir}^* \Xi_{js} \int_t^{t+T} v_r^*(t')v_s(t')dt' = \sum_{r,s} \Xi_{ir}^* \Xi_{js} \Omega_{rs}(t;T) \tag{41}$$

Here we are making use of the spectral decomposition in Eq. (7) and we define $\Omega_{rs}(t;T)$ as the estimated covariance between the OU modes v_r, v_s. From the simulations, it results that only a few slow uncorrelated modes have a significant weight Ξ in the decomposition. In this scenario, the covariance matrix is well approximated by the dot product between the representational vectors in the functional space:

$$\mathrm{cov}_{ij}(t;T) \approx \boldsymbol{\Xi}_i \cdot \boldsymbol{\Xi}_j \tag{42}$$

Consequently for the correlation matrix it holds

$$\mathrm{FC}_{ij}(t;T) \approx \mathrm{S_C}\left(\boldsymbol{\Xi}_i, \boldsymbol{\Xi}_j\right) \tag{43}$$

where $\mathrm{S_C}$ indicates the cosine similarity function, i.e., the Pearson correlation between row and column vectors composing $\boldsymbol{\Xi}^*$ and $\boldsymbol{\Xi}$, respectively, under the assumption of a zero mean for the system observables.

Simulation parameters

Network parameters have been selected through grid search methods to optimize prediction accuracy. The reservoir state is initialized in a zero-activity state and is disrupted from equilibrium by the input for a transient period of $20\,s$ before learning. This period corresponds to the t^* mentioned in the preceding paragraphs. The input signal is linearly interpolated, and the system's dynamics are resolved using the Euler–Maruyama method with a time step that is one-hundredth of the sampling interval. The spectral radius ρ of the synaptic matrix, used for fMRI analysis, is set to 0.5, and the time constant τ is 1.0. The standard deviation g^{in} of the input matrix $\boldsymbol{W}^{\mathrm{in}}$ is set to 1. The thermal noise shown in Fig. 3 has a regularizing effect. Consequently, we do not employ further regularization for the readout regression. For the example system depicted in Fig. 1, the parameters are $\rho = 0.5$, $\tau = 0.1$ with no regularizers, and g^{in} equal to the inverse of the square root of the number of inputs. The stochastic scenario is characterized by thermal noise with a standard deviation of 10^{-5} and ridge regression with a regularization of 10^{-7} . For cases with missing variables, the parameters are set to $\rho = 0.6$, $\tau = 0.25$.

Clustering analysis

The clustering analysis is performed on the resulting spectrum of the data by discretizing the complex space into a grid with a real resolution of $0.00025s^{-1}$, and an imaginary resolution of $0.01s^{-1}$.

To further reduce the dimensionality of the data and get an interpretable 3 dimensional representation, Linear Discriminant Analysis (LDA) is employed. LDA is a widely used technique in machine learning (69). It is a supervised learning method that aims to reduce the dimensionality of the data while preserving the discriminatory information between different classes. The core idea is to find a linear transformation that maximizes the ratio of the between-class scatter $\boldsymbol{S}_b$ to the within-class scatter $\boldsymbol{S}_w$. This optimization process equals to solve the following generalized eigenvalue problem

$$\boldsymbol{S}_b \boldsymbol{l} = \lambda \boldsymbol{S}_w \boldsymbol{l} \tag{44}$$

The optimal projection matrix $\boldsymbol{L}$ to a subspace of dimension k is given by the eigenvectors $\boldsymbol{l}$ corresponding to the largest k eigenvalues.

Linear Decoders

To find the distinct frequency patterns that characterize each brain area, we computed the associated linear decoder within the discretized frequency-vector representation. This process involved linking each brain area to a target versor in an 11-dimensional space. Then we employed a pseudoinverse algorithm to derive the decoders, which maps the frequency-vectors – resulting from all subjects and 100 reservoir realizations – to their corresponding relative versors.

Manuscript Language Optimization

To enhance the clarity, coherence and readability of the manuscript, we employed large language models (LLMs) for language editing and refinement. All edits assisted by the model were subsequently reviewed by us to ensure accuracy and to preserve the intended scientific meaning.

References

[1] Christopher M. Bishop. *Pattern recognition and machine learning*. Springer, New York, NY, 2006.

[2] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. `http://www.deeplearningbook.org`.

[3] Bernhard Mehlig. *Machine learning with neural networks: an introduction for scientists and engineers*. Cambridge University Press, 2021.

[4] Herbert Jaeger and Harald Haas. Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science*, 304(5667):78–80, 2004.

[5] Mantas Lukoševičius and Herbert Jaeger. Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.*, 3(3):127–149, 2009.

[6] Jaideep Pathak, Brian Hunt, Michelle Girvan, Zhixin Lu, and Edward Ott. Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Phys. Rev. Lett.*, 120(2): 024102, 2018.

[7] Pantelis-Rafail Vlachas, Jaideep Pathak, Brian R Hunt, Themistoklis P Sapsis, Michelle Girvan, Edward Ott, and Petros Koumoutsakos. Backpropagation algorithms and reservoir computing in recurrent neural networks for the forecasting of complex spatiotemporal dynamics. *Neural Networks*, 126:191–217, 2020.

[8] Keshav Srinivasan, Nolan Coble, Joy Hamlin, Thomas Antonsen, Edward Ott, and Michelle Girvan. Parallel machine learning for forecasting the dynamics of complex networks. *Phys. Rev. Lett.*, 128(16):0–4, 2022. doi: $10.1103/\mathrm{PhysRevLett.128.164101}$.

[9] H B Barlow. Possible principles underlying the transformation of sensory messages. In W A Rosenblith, editor, *Sensory Communication*, pages 217–34. MIT Press, Cambridge, MA, 1961.

[10] Eero P Simoncelli and Bruno A Olshausen. Natural image statistics and neural representation. *Annu. Rev. Neurosci.*, 24(1):1193–1216, 2001. doi: 10.1146/annurev.neuro.24.1.1193.

[11] Stefano Fusi, Earl K Miller, and Mattia Rigotti. Why neurons mix: high dimensionality for higher cognition. *Curr. Opin. Neurobiol.*, 37:66–74, 2016. doi: 10.1016/j.conb.2016.01.010.

[12] Valerio Mante, David Sussillo, Krishna V Shenoy, and William T Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84, 2013.

[13] David Sussillo and O. Barak. Opening the black box: Low-dimensional dynamics in high-dimensional recurrent neural netw. *Neural Computation*, 25:626–649, 2013. doi: 10.1162/NECO_a_00409.

[14] Daniel Durstewitz, Georgia Koppe, and Max Ingo Thurm. Reconstructing computational system dynamics from neural data with recurrent neural networks. *Nat. Rev. Neurosci.*, 24(11):693–710, 2023.

[15] Gabriele Di Antonio, Sofia Raglio, and Maurizio Mattia. A geometrical solution underlies general neural principle for serial ordering. *Nat. Commun.*, 15(1):8238, 2024.

[16] Christoph Molnar, Giuseppe Casalicchio, and Bernd Bischl. Interpretable machine learning–a brief history, state-of-the-art and challenges. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 417–431. Springer, 2020.

[17] William Gilpin. Generative learning for nonlinear dynamics. *Nat. Rev. Phys.*, 6(3):194–206, 2024. doi: 10.1038/s42254-024-00688-2.

[18] Hugues Turbé, Mina Bjelogrlic, Christian Lovis, and Gianmarco Mengaldo. Evaluation of post-hoc interpretability methods in time-series classification. *Nat. Mach. Intell.*, 5(3):250–260, 2023. doi: 10.1038/s42256-023-00620-w.

[19] Alexandre Mauroy, Y Susuki, and Igor Mezić. *Koopman operator in systems and control*. Springer, 2020.

[20] Steven L. Brunton, Marko Budišić, Eurika Kaiser, and J. Nathan Kutz. Modern Koopman theory for dynamical systems. *SIAM Rev.*, 64(2):229–340, 2022. doi: 10.1137/21M1401243.

[21] Peter J Schmid. Dynamic Mode Decomposition and its variants. *Annu. Rev. Fluid Mech.*, 54(1):225–254, 2022. doi: 10.1146/annurev-fluid-030121-015835.

[22] Gabriele Di Antonio and Gianni Valerio Vinci. Koopman correlations underlie linear response and causality. *arXiv preprint arXiv:2410.08708*, 2024.

[23] Fei Tao, Bin Xiao, Qinglin Qi, Jiangfeng Cheng, and Ping Ji. Digital twin modeling. *J. Manuf. Syst.*, 64:372–389, 2022. doi: 10.1016/j.jmsy.2022.06.015.

[24] Steven A Niederer, Michael S Sacks, Mark Girolami, and Karen Willcox. Scaling digital twins from the artisanal to the industrial. *Nat. Comput. Sci.*, 1(5):313–320, 2021. doi: 10.1038/s43588-021-00072-5.

[25] Michael Schirner, Simon Rothmeier, Viktor K Jirsa, Anthony Randal McIntosh, and Petra Ritter. An automated pipeline for constructing personalized virtual brains from multimodal neuroimaging data. *NeuroImage*, 117:343–357, 2015.

[26] Viktor K Jirsa, Timothée Proix, D Perdikis, MM Woodman, H Wang, J Gonzalez-Martinez, Christophe Bernard, C Bénar, Maxime Guye, Patrick Chauvel, and Fabrice Bartolomei. The Virtual Epileptic Patient: Individualized whole-brain models of epilepsy spread. *NeuroImage*, 145:377–388, 2017. doi: 10.1016/j.neuroimage.2016.04.049.

[27] Viktor Jirsa, Huifang Wang, Paul Triebkorn, Meysam Hashemi, Jayant Jha, Jorge Gonzalez-Martinez, Maxime Guye, Julia Makhalova, and Fabrice Bartolomei. Personalised virtual brain models in epilepsy. *Lancet Neurol.*, 22(5):443–454, 2023. doi: 10.1016/S1474-4422(23)00008-X.

[28] Lucius Samo Fekonja, Robert Schenk, Emily Schröder, Rosario Tomasello, Samo Tomšič, and Thomas Picht. The digital twin in neuroscience: from theory to tailored therapy. *Front. Neurosci.*, 18:1–8, 2024. doi: 10.3389/fnins.2024.1454856.

[29] Ljupco Kocarev and U Parlitz. Generalized synchronization, predictability, and equivalence of unidirectionally coupled dynamical systems. *Phys. Rev. Lett.*, 76(11):1816–9, 1996. doi: 10.1103/PhysRevLett.76.1816.

[30] Zhixin Lu, Brian R Hunt, and Edward Ott. Attractor reconstruction by machine learning. *Chaos*, 28(6), 2018. doi: 10.1063/1.5039508.

[31] Herbert Jaeger. The "echo state" approach to analysing and training recurrent neural networks. Technical report, German National Research Center for Information Technology, 2001.

[32] Izzet B Yildiz, Herbert Jaeger, and Stefan J Kiebel. Re-visiting the echo state property. *Neural Netw.*, 35: 1–9, 2012. doi: 10.1016/j.neunet.2012.07.005.

[33] Floris Takens. Detecting strange attractors in turbulence. In David Rand and Lai-Sang Young, editors, *Dynamical Systems and Turbulence, Warwick 1980*, pages 366–381, Berlin, Heidelberg, 1981. Springer Berlin Heidelberg. ISBN 9783540389453.

[34] Tim Sauer, James A Yorke, and Martin Casdagli. Embedology. *J. Stat. Phys.*, 65(3-4):579–616, 1991. doi: 10.1007/BF01053745.

[35] J Stark, D S Broomhead, M E Davies, and J Huke. Delay embeddings for forced systems. II. Stochastic forcing. *J. Nonlinear Sci.*, 13:519–577, 2003. doi: 10.1007/s00332-003-0534-4.

[36] Erik M Bollt. On explaining the surprising success of reservoir computing forecaster of chaos? the universal machine learning dynamical system with contrast to var and dmd. *Chaos*, 31(1), 2021.

[37] Daniel J Gauthier, Erik M Bollt, Aaron Griffith, and Wendson AS Barbosa. Next generation reservoir computing. *Nat. Commun.*, 12(1):1–8, 2021.

[38] Pietro Verzelli, Cesare Alippi, Lorenzo Livi, and Peter Tiňo. Input-to-state representation in linear reservoirs dynamics. *IEEE Trans. Neural Netw. Learn. Syst.*, 33(9):4598–4609, 2021.

[39] Zheng-Meng Zhai, Ling-Wei Kong, and Ying-Cheng Lai. Emergence of a resonance in machine learning. *Phys. Rev. Res.*, 5(3):033127, 2023.

[40] E. J W Van Someren, Y. D. Van Der Werf, Pieter R Roelfsema, H. D. Mansvelder, and F. H. Lopes da Silva. Slow brain oscillations of sleep, resting state, and vigilance. *Prog. Brain Res.*, 193:3–15, 2011. doi: 10.1016/B978-0-444-53839-0.00001-6.

[41] Sandro Lecci, Laura M J Fernandez, Frederik D Weber, Romain Cardis, Jean Yves Chatton, Jan Born, and Anita Lüthi. Coordinated infraslow neural and cardiac oscillations mark fragility and offline periods in mammalian sleep. *Sci. Adv.*, 3(2), 2017. doi: 10.1126/sciadv.1602026.

[42] Daniel Gutierrez-Barragan, M Albert Basson, Stefano Panzeri, and Alessandro Gozzi. Infraslow state fluctuations govern spontaneous fMRI network dynamics. *Curr. Biol.*, 29(14):2295–2306.e5, 2019. doi: 10.1016/j.cub.2019.06.017.

[43] Núria Tort-Colet, Cristiano Capone, Maria V Sanchez-Vives, and Maurizio Mattia. Attractor competition enriches cortical dynamics during awakening from anesthesia. *Cell Rep.*, 35(12):109270, 2021. doi: 10.1016/j.celrep.2021.109270.

[44] Erfan Nozari, Maxwell A Bertolero, Jennifer Stiso, Lorenzo Caciagli, Eli J Cornblath, Xiaosong He, Arun S Mahadevan, George J Pappas, and Dani S Bassett. Macroscopic resting-state brain dynamics are best described by linear models. *Nat. Biomed. Eng.*, 8(1):68–84, 2024. doi: 10.1038/s41551-023-01117-y.

[45] Holger Kantz and Thomas Schreiber. *Nonlinear time series analysis*. Cambridge university press, 2003.

[46] Hassan Arbabi and Igor Mezic. Ergodic theory, dynamic mode decomposition, and computation of spectral properties of the koopman operator. *SIAM J. Appl. Dyn. Syst.*, 16(4):2096–2126, 2017.

[47] Steven L Brunton, Bingni W Brunton, Joshua L Proctor, Eurika Kaiser, and J Nathan Kutz. Chaos as an intermittently forced linear system. *Nat. Commun.*, 8(1):19, 2017.

[48] Ru Zhang, Leon Aksman, Dilmini Wijesinghe, John M Ringman, Danny J J Wang, and Kay Jann. A longitudinal study of functional brain complexity in progressive Alzheimer's disease. *Alzheimer's and Dementia*, 17 (1):1–11, 2025. doi: 10.1002/dad2.70059.

[49] David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy E J Behrens, Essa Yacoub, Kamil Ugurbil, and WU-Minn HCP Consortium. The WU-Minn Human Connectome Project: An overview. *NeuroImage*, 80: 62–79, 2013. doi: 10.1016/j.neuroimage.2013.05.041.

[50] Michael P Harms, Leah H Somerville, Beau M Ances, Jesper Andersson, Deanna M Barch, Matteo Bastiani, Susan Y Bookheimer, Timothy B Brown, Randy L Buckner, Gregory C Burgess, Timothy S Coalson, Michael A Chappell, Mirella Dapretto, Gwenaëlle Douaud, Bruce Fischl, Matthew F Glasser, Douglas N Greve, Cynthia Hodge, Keith W Jamison, Saad Jbabdi, Sridhar Kandala, Xiufeng Li, Ross W Mair, Silvia Mangia, Daniel Marcus, Daniele Mascali, Steen Moeller, Thomas E Nichols, Emma C Robinson, David H Salat, Stephen M Smith, Stamatios N Sotiropoulos, Melissa Terpstra, Kathleen M Thomas, M Dylan Tisdall, Kamil Ugurbil, Andre van der Kouwe, Roger P Woods, Lilla Zöllei, David C Van Essen, and Essa Yacoub. Extending the Human Connectome Project across ages: Imaging protocols for the Lifespan Development and Aging projects. *NeuroImage*, 183:972–984, 2018. doi: 10.1016/j.neuroimage.2018.09.060.

[51] Susan Y Bookheimer, David H Salat, Melissa Terpstra, Beau M Ances, Deanna M Barch, Randy L Buckner, Gregory C Burgess, Sandra W Curtiss, Mirella Diaz-Santos, Jennifer Stine Elam, Bruce Fischl, Douglas N Greve, Hannah A Hagy, Michael P Harms, Olivia M Hatch, Trey Hedden, Cynthia Hodge, Kevin C Japardi, Taylor P Kuhn, Timothy K Ly, Stephen M Smith, Leah H Somerville, Kâmil Uğurbil, Andre van der Kouwe, David Van Essen, Roger P Woods, and Essa Yacoub. The Lifespan Human Connectome Project in Aging: An overview. *NeuroImage*, 185:335–348, 2019. doi: 10.1016/j.neuroimage.2018.10.009.

[52] Yonatan Sanz Perl, Sol Fittipaldi, Cecilia Gonzalez Campo, Sebastián Moguilner, Josephine Cruzat, Matias E Fraile-Vazquez, Rubén Herzog, Morten L Kringelbach, Gustavo Deco, Pavel Prado, Agustin Ibanez, and Enzo Tagliazucchi. Model-based whole-brain perturbational landscape of neurodegenerative diseases. *eLife*, 12:1–25, 2023. doi: 10.7554/eLife.83970.

[53] Lucine L Oganesian and Maryam M Shanechi. Brain–computer interfaces for neuropsychiatric disorders. *Nat. Rev. Bioeng.*, 2(8):653–670, 2024.

[54] Ling-Wei Kong, Gene A Brewer, and Ying-Cheng Lai. Reservoir-computing based associative memory and itinerancy for complex dynamical attractors. *Nat. Commun.*, 15(1):4840, 2024. doi: 10.1038/s41467-024-49190-4.

[55] Hillel Sanhedrai, Jianxi Gao, Amir Bashan, Moshe Schwartz, Shlomo Havlin, and Baruch Barzel. Reviving a failed network through microscopic interventions. *Nat. Phys.*, 18(3):338–349, 2022. doi: 10.1038/s41567-021-01474-y.

[56] Gustavo Deco and Viktor K Jirsa. Ongoing cortical activity at rest: criticality, multistability, and ghost attractors. *J. Neurosci.*, 32(10):3366–75, 2012. doi: 10.1523/JNEUROSCI.2523-11.2012.

[57] Maria V Sanchez-Vives, Marcello Massimini, and Maurizio Mattia. Shaping the default activity pattern of the cortical network. *Neuron*, 94(5):993–1001, 2017. doi: 10.1016/j.neuron.2017.05.015.

[58] Serena di Santo, Pablo Villegas, Raffaella Burioni, and Miguel A Muñoz. Landau-Ginzburg theory of cortex dynamics: Scale-free avalanches emerge at the edge of synchronization. *Proc. Natl. Acad. Sci. USA*, 115 (7):E1356–E1365, 2018. doi: 10.1073/pnas.1712989115.

[59] Chen Song, Melanie Boly, Enzo Tagliazucchi, Helmut Laufs, and Giulio Tononi. fMRI spectral signatures of sleep. *Proc. Natl. Acad. Sci. USA*, 119(30):1–12, 2022. doi: 10.1073/pnas.2016732119.

[60] Adam J Eisen, Leo Kozachkov, André M Bastos, Jacob A Donoghue, Meredith K Mahnke, Scott L Brincat, Sarthak Chandra, John Tauber, Emery N Brown, Ila R Fiete, and Earl K Miller. Propofol anesthesia destabilizes neural dynamics across cortex. *Neuron*, 112(16):2799–2813.e9, 2024. doi: 10.1016/j.neuron.2024.06.011.

[61] Georgia Koppe, Hazem Toutounji, Peter Kirsch, Stefanie Lis, and Daniel Durstewitz. Identifying nonlinear dynamical systems via generative recurrent neural networks with applications to fMRI. *PLOS Comput. Biol.*, 15(8):e1007263, 2019. doi: 10.1371/journal.pcbi.1007263.

[62] Matthieu Gilson, Enzo Tagliazucchi, and Rodrigo Cofré. Entropy production of multivariate Ornstein-Uhlenbeck processes correlates with consciousness levels in the human brain. *Phys. Rev. E*, 107(2):1–6, 2023. doi: 10.1103/PhysRevE.107.024121.

[63] Viktor Sip, Meysam Hashemi, Timo Dickscheid, Katrin Amunts, Spase Petkoski, and Viktor Jirsa. Characterization of regional differences in resting-state fMRI with a data-driven network model of brain dynamics. *Sci. Adv.*, 9(11), 2023. doi: 10.1126/sciadv.abq7547.

[64] Jonathan D Touboul and Alain Destexhe. Can power-law scaling and neuronal avalanches arise from stochastic dynamics? *PLoS ONE*, 5(2), 2010. doi: 10.1371/journal.pone.0008982.

[65] Andrea Santoro, Federico Battiston, Giovanni Petri, and Enrico Amico. Higher-order organization of multivariate time series. *Nat. Phys.*, 19(2):221–229, 2023. doi: 10.1038/s41567-022-01852-0.

[66] Gustavo Deco and Morten L Kringelbach. Turbulent-like Dynamics in the Human Brain. *Cell Rep.*, 33(10): 108471, 2020. doi: 10.1016/j.celrep.2020.108471.

[67] Mikail Khona and Ila R Fiete. Attractor and integrator networks in the brain. *Nat. Rev. Neurosci.*, 23, 2022. doi: 10.1038/s41583-022-00642-0.

[68] Braden A W Brinkman, Han Yan, Arianna Maffei, Il Memming Park, Alfredo Fontanini, Jin Wang, and Giancarlo La Camera. Metastable dynamics of neural circuits and networks. *Appl. Phys. Rev.*, 9(1):011313, 2022. doi: 10.1063/5.0062603.

[69] Richard O Duda, Peter E Hart, and David D Stork. *Pattern classification*. John Wiley & Sons, New York, NY, 2nd edition edition, 2001.

[70] Bernard O Koopman. Hamiltonian systems and transformation in hilbert space. *Proc. Natl. Acad. Sci. U.S.A.*, 17(5):315–318, 1931.

[71] Alessandro Treves. Mean-field analysis of neuronal spike dynamics. *Network*, 4(3):259–84, 1993.

[72] Bruce W Knight. Dynamics of encoding in neuron populations: some general mathematical features. *Neural Comput.*, 12(3):473–518, 2000.

[73] Maurizio Mattia and Paolo Del Giudice. Population dynamics of interacting spiking neurons. *Phys. Rev. E*, 66(5 Pt 1):051917, 2002. doi: 10.1103/PhysRevE.66.051917.

[74] Edward T Bullmore and Olaf Sporns. The economy of brain network organization. *Nat. Rev. Neurosci.*, 13 (5):336–349, 2012. doi: 10.1038/nrn3214.

[75] Enrique C A Hansen, Demian Battaglia, Andreas Spiegler, Gustavo Deco, and Viktor K Jirsa. Functional connectivity dynamics: Modeling the switching behavior of the resting state. *NeuroImage*, 105:525–535, 2015. doi: 10.1016/j.neuroimage.2014.11.001.

[76] Gustavo Deco, Morten L Kringelbach, Viktor K Jirsa, and Petra Ritter. The dynamics of resting fluctuations in the brain: Metastability and its dynamical cortical core. *Sci. Rep.*, 7(1):1–14, 2017. doi: 10.1038/s41598-017-03073-5.

Acknowledgments

We thank M. Allegra for the feedback on an earlier version of the manuscript and S. P. Caminiti, I. Branchi and G. V. Vinci for the stimulating discussions. Work partially funded by the NEXTGENERATIONEU and MUR (PNRR-M4C2I1.3) project MNESYS (PE0000006–DD 1553 11.10.2022) and project EBRAINS-Italy (IR0000011– DD 101 16.6.2022) to MM.