Functional Connectivity-based Attractor Dynamics in Rest, Task, and Disease

Robert Englert[1,2], Balint Kincses[1,3], Raviteja Kotikalapudi[1,3], Giuseppe Gallitto[1,3], Jialin Li[1,3,4], Kevin Hoffschlag[1,3], Choong-Wan Woo[5,6], Tor D. Wager[7], Dagmar Timmann[1,3], Ulrike Bingel[1,3], Tamas Spisak[1,2]*

1 Center for Translational Neuro- and Behavioral Sciences (C-TNBS), University Medicine Essen, Germany
2 Department of Diagnostic and Interventional Radiology and Neuroradiology, University Medicine Essen, Germany
3 Department of Neurology, University Medicine Essen, Germany
4 Max Planck School of Cognition, Leipzig, Germany
5 Center for Neuroscience Imaging Research, Institute for Basic Science, Suwon, South Korea
6 Department of Biomedical Engineering, Sungkyunkwan University, Suwon, South Korea
7 Department of Psychological and Brain Sciences, Dartmouth College, Hanover, NH, USA
** corresponding author: tamas.spisak@uk-essen.de*

Key Points:

- We present a simple yet powerful generative computational model for large-scale brain dynamics
- Based on the theory of artificial attractor neural networks emerging from first principles of self-organization
- Model dynamics accurately reconstruct several characteristics of resting-state brain dynamics and confirm theoretical predictions of emergent attractor self-orthogonalization
- Our model captures both task-induced and pathological changes in brain activity
- fcANNs offer a simple and interpretable computational alternative to conventional descriptive analyses of brain function

Project website (with interactive manuscript): https://pni-lab.github.io/connattractor

Abstract

Functional brain connectivity has been instrumental in uncovering the large-scale organization of the brain and its relation to various behavioral and clinical phenotypes. Understanding how this functional architecture relates to the brain's dynamic activity repertoire is an essential next step towards interpretable generative models of brain function. We propose functional connectivity-based Attractor Neural Networks (fcANNs), a theoretically inspired model of macro-scale brain dynamics, simulating recurrent activity flow among brain regions based on first principles of self-organization. In the fcANN framework, brain dynamics are understood in relation to attractor states; neurobiologically meaningful activity configurations that minimize the free energy of the system. We provide the first evidence that large-scale brain attractors - as reconstructed by fcANNs - exhibit an approximately orthogonal organization, which is a signature of the self-orthogonalization mechanism of the underlying theoretical framework of free-energy-minimizing attractor networks. Analyses of 7 distinct datasets demonstrate that fcANNs can accurately reconstruct and predict brain dynamics under a wide range of conditions, including resting and task states, and brain disorders. By establishing a formal link between connectivity and activity, fcANNs offer a simple and interpretable computational alternative to conventional descriptive analyses.

Introduction

Brain function is characterized by the continuous activation and deactivation of anatomically distributed neuronal populations (Buzsaki, 2006). Irrespective of the presence or absence of explicit stimuli, brain regions appear to work in concert, giving rise to rich and spatiotemporally complex fluctuations (Bassett & Sporns, 2017). These fluctuations are not random (Liu & Duyn, 2013; Zalesky *et al.*, 2014); they organize around large-scale gradients (Margulies *et al.*, 2016; Huntenburg *et al.*, 2018) and exhibit quasi-periodic properties, with a limited number of recurring patterns often termed as "brain substates" (Greene *et al.*, 2023; Kringelbach & Deco, 2020; Vidaurre *et al.*, 2017; Liu & Duyn, 2013). A wide variety of descriptive techniques have been previously employed to characterize whole-brain dynamics (Smith *et al.*, 2012; Vidaurre *et al.*, 2017; Liu & Duyn, 2013; Chen *et al.*, 2018). These efforts have provided accumulating evidence not only for the existence of dynamic brain substates but also for their clinical significance (Hutchison *et al.*, 2013; Barttfeld *et al.*, 2015; Meer *et al.*, 2020). However, the underlying driving forces remain elusive due to the descriptive nature of such studies.

Conventional computational approaches attempt to solve this puzzle by going all the way down to the biophysical properties of single neurons, and aim to construct a model of larger neural populations, or even the entire brain (Breakspear, 2017). These approaches have shown numerous successful applications (Murray *et al.*, 2018; Kriegeskorte & Douglas, 2018; Heinz *et al.*, 2019). However, such models need to estimate a vast number of neurobiologically motivated free parameters to fit the data. This hampers their ability to effectively bridge the gap between explanations at the level of single neurons and the complexity of behavior (Breakspear, 2017). Recent efforts using coarse-grained brain network models (Schirner *et al.*, 2022; Schiff *et al.*, 1994; Papadopoulos *et al.*, 2017; Seguin *et al.*, 2023) and linear network control theory (Chiêm *et al.*, 2021; Scheid *et al.*, 2021; Gu *et al.*, 2015) opted to trade biophysical fidelity to phenomenological validity. Such models have provided insights into some of the inherent key characteristics of the brain as a dynamic system; for instance, the importance of stable patterns, the *attractor states*, in governing brain dynamics (Deco *et al.*, 2012; Golos *et al.*, 2015; Hansen *et al.*, 2015). While attractor networks become established models of micro-scale canonical brain circuits in the last four decades (Khona & Fiete, 2022), these studies suggest that attractor dynamics are essential characteristics of macro-scale brain dynamics as well (Poerio & Karapanagiotidis, 2025). Attractor networks, however, come in many flavors and the specific forms and behaviors of these networks are heavily influenced by the chosen inference and learning rules, making it unclear which variety should be in focus when modeling brain dynamics. Given that the brain showcases not only multiple signatures of attractor dynamics but also the ability to evolve and adapt through self-organization (i.e., in the absence of any centralized control), investigating attractor models from the point of view of self-organization may be key to narrow down the set of viable models.

In our recent theoretical work (Spisak & Friston, 2025), we identified the class of attractor networks that emerge from first principles of self-organization - as articulated by the Free Energy Principle (FEP) (Friston, 2010; Friston *et al.*, 2023) - and identified the emergent inference and learning rules guiding the dynamics of such systems. This theoretical framework reveals that the minimization of variational free energy locally - e.g., by individual network nodes - gives rise to a dual dynamic: simultaneous inference (updating activity) and learning (optimizing connectivity). The emergent inference process in these systems is equivalent to local Bayesian update dynamics for the individual network nodes, homologous to the stochastic relaxation observed in conventional Boltzmann neural network architectures (e.g. stochastic Hopfield networks, (Hopfield, 1982; Koiran, 1994)), and in line with the empirical observation that activity in the brain "flows" following similar dynamics (Cole *et al.*, 2016; Sanchez-Romero *et al.*, 2023; Cole, 2024). Importantly, in this framework, attractor states are not simply an epiphenomenon of collective dynamics, but serve as global priors in the Bayesian sense, that get combined with the current activity configuration so that the updated activity samples from the posterior (akin to a Markov-Chain Monte Carlo (MCMC) sampling process).

Learning, on the other hand, emerges in this framework in the form of a distinctive coupling plasticity - a local, incremental learning rule - that continuously adjusts coupling weights to preserve low free energy in anticipation of future sensory encounters following a contrastive predictive coding scheme (Millidge *et al.*, 2022), effectively implementing action selection in the active inference sense (Friston *et al.*, 2016). Importantly, the learning dynamics emerging in our theoretical framework provide a strong, testable hypothesis: if the brain operates as a free energy minimizing attractor network, its large-scale attractors should be approximately orthogonal to each other. This is not a general property of all recurrent (attractor) neural networks, but a direct consequence of free energy minimization, shown both mathematically and with simulations in (Spisak & Friston, 2025). As our theoretical framework - by design - embraces multiple valid levels of description (through coarse-graining), it is well-suited to serve as a basis for a computational model of large-scale brain dynamics.

In the present work, we translate the results of this novel theoretical framework into a computational model of macro-scale brain dynamics and deploy a diverse set of experimental, clinical, and meta-analytic studies to perform an initial investigation of several of its predictions. We start by showing that - if large-scale brain

dynamics evolve and organize according to the emergent rules of this framework - the corresponding attractor model can be effectively approximated from functional connectivity data, as measured with resting-state fMRI. Based on the network topology spanned by functional connectivity, our model assigns a free energy level for any arbitrary activation pattern and determines a "trajectory of least action" towards one of a finite number of attractor states that minimize this energy (Figure 1). We then perform an initial investigation of the robustness and biological plausibility of the attractor states of the reconstructed network and whether it is able to reproduce various characteristics of resting-state brain dynamics. Importantly, we directly test the framework's prediction on the emergence of (approximately) orthogonal attractor states. Capitalizing on the generative nature of our model, we also demonstrate how it can capture - and potentially explain - the effects of various perturbations and alterations of these dynamics, from task-induced activity to changes related to brain disorders.

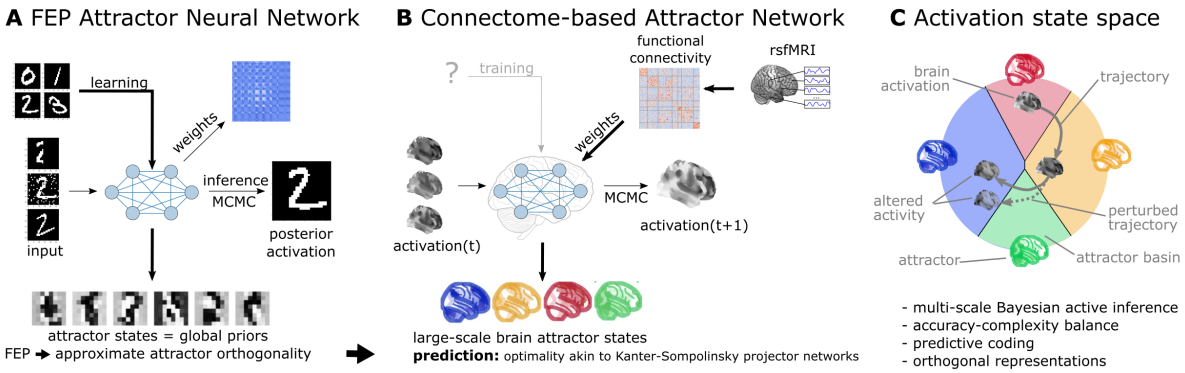

Figure 1: Functional connectivity-based attractor neural networks as models of macro-scale brain dynamics.
A Free-energy-minimizing artificial neural networks (Spisak & Friston, 2025) are a form of recurrent stochastic artificial neural networks that, similarly to classical Hopfield networks (Hopfield, 1982; Koiran, 1994), can serve as content-addressable ("associative") memory systems. More generally, through the learning rule emerging from local free-energy minimization, the weights of these networks will encode a global internal model of the external world. The priors of this internal generative model are represented by the attractor states of the network that - as a special consequence of free-energy minimization - will tend to be orthogonal to each other. During stochastic inference (local free-energy minimization), the network samples from the posterior that combines these priors with the previous brain substates (also encompassing incoming stimuli), akin to Markov chain Monte Carlo (MCMC) sampling. **B** In accordance with this theoretical framework, we consider regions of the brain as nodes of a free-energy-minimizing artificial neural network. Instead of initializing the network with the structural wiring of the brain or training it to solve specific tasks, we set its weights empirically, using information about the interregional "activity flow" across regions, as estimated via functional brain connectivity. Applying the inference rule of our framework - which displays strong analogies with the relaxation rule of Hopfield networks and the activity flow principle that links activity to connectivity in brain networks - results in a generative computational model of macro-scale brain dynamics, that we term a functional connectivity-based (stochastic) attractor neural network (fcANN). **C** The proposed computational framework assigns a free energy level, a probability density and a "trajectory of least action" towards an attractor state to any brain activation pattern and predicts changes of the corresponding dynamics in response to alterations in activity and/or connectivity. The theoretical framework underlying the fcANNs - based on the assumption that the brain operates as a free energy minimizing attractor network - draws formal links between attractor dynamics and multi-level Bayesian active inference.

Theoretical background

Free energy minimizing artificial neural networks

The computational model at the heart of this work is a direct implementation of the inference dynamics that emerge from a recently proposed theoretical framework of free energy minimizing (self-orthogonalizing) attractor networks (Spisak & Friston, 2025) (FEP-ANNs). FEP-ANNs are a class of attractor neural networks that emerge from the Free Energy Principle (FEP) (Friston, 2010; Friston *et al.*, 2023). The FEP posits that any self-organizing random dynamical system that maintains its integrity over time (i.e. has a steady-state distribution and a statistical separation from its environment) must act in a way that minimizes its variational free energy. FEP-ANNs apply the FEP recursively. We assume a network of N units, where each unit is represented by a single continuous-valued state variable $\sigma_i \in [-1,1]$, so that the activity of the network is described by a vector of states $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ and these states are conditionally independent of each other, given boundary states that realize the necessary statistical separation between them (corresponding to a complex Markov blanket structure in the

FEP terminology). When assuming σ_i states that follow a continuous Bernoulli (a.k.a. truncated exponential: $p(\sigma_i) \propto e^{\kappa_i \sigma_i}$) distribution (parameterized by the single parameter κ_i) and deterministic couplings J, the steady-state distribution can be expressed as:

$$p^*(\sigma) \propto \exp\left(\sum_i b_i \sigma_i + \frac{1}{2}\sum_{i,j} J^S_{ij} \sigma_i \sigma_j\right) \tag{1}$$

where b_i represents the local evidence or bias for each unit i (e.g. external input or intrinsic excitability of a brain region), $J^S = \frac{1}{2}(J + J^T)$ is the symmetric component of the coupling weights between units i and j, and β is an inverse temperature or precision parameter. Note that while this steady-state distribution has the same functional form as continuous-state Boltzmann machines or stochastic Hopfield networks, the true coupling weights J do not have to be symmetric as usually assumed in those architectures. Asymmetric couplings break detailed balance, meaning that p^* is no longer an equilibrium distribution. However, the antisymmetric component $J^A = \frac{1}{2}(J - J^T)$ does not contribute to the steady-state distribution p^* as it only induces circulating (solenoidal) flow in the state space which is tangential to the level sets of p^*. Thus, while the overall framework can describe general attractor networks with asymmetric couplings and non-equilibrium steady states (NESS), it also implies that knowing only the symmetric component of the coupling weights is sufficient to reconstruct the steady-state distribution p^* of the underlying system. This is a highly useful property for the purposes of the present study, where the couplings are reconstructed from resting state fMRI data, without any explicit information about the directionality of functional connections. For a detailed derivation of the steady state distribution, see (Spisak & Friston, 2025) and Supplementary Information 1.

Knowing the steady-state distribution of a free-energy-minimizing attractor network, we can derive two types of emergent dynamics from the single imperative of free energy minimization: inference and learning.

Inference: stochastic relaxation dynamics

Inference arises from minimizing free energy with respect to the states σ. For a single unit, this yields a local update rule homologous to the relaxation dynamics in Hopfield networks:

$$E_q[\sigma_i] = L(b_q) = \underbrace{L\left(\underbrace{b_i}_{bias} + \underbrace{\sum_{j\neq i} J_{ij}\sigma_j}_{local\ potential}\right)}_{sigmoid\ (Langevin)} \tag{2}$$

where L is a sigmoidal activation function (a Langevin function in our case). This rule dictates that each unit updates its activity stochastically, based on a weighted sum of the activity of other units, plus its own intrinsic bias. See (Spisak & Friston, 2025) and Supplementary Information 2 for a detailed derivation of the inference dynamics.

Note that the rule is expressed in terms of the expected value of the state σ_i, which is a stochastic quantity. However, in the limiting case of symmetric couplings (which is the case throughout the present study) and least-action dynamics (i.e. no noise), this update rule reduces to the classical relaxation dynamics of (continuous-state) Hopfield networks. In the present study, we use both the deterministic ("least action") and stochastic variants of the inference rule. The former identifies attractor states; the latter serves as a generative model for large-scale, multistable brain dynamics.

In the present study, we make the simplifying assumption that all nodes have zero bias ($b = 0$). Furthermore, we allow investigating different scaling factors for the J couplings matrix (given the uncertainties around the magnitude of association in the functional connectome) by introducing a "scaling factor" β. This leads to the following update rule:

$$\sigma_i^{(t+1)} = L\left(\beta \sum_{j\neq i} J_{ij}\, \sigma_j^{(t)}\right) + \text{noise} \tag{3}$$

The scaling factor β is analogous to the inverse temperature parameter known in Hopfield networks and Boltzmann machines.

In the basis framework (Spisak & Friston, 2025), inference is a gradient descent on the variational free energy landscape with respect to the states σ and can be interpreted as a form of approximate Bayesian inference, where the expected value of the state σ_i is interpreted as the posterior mean given the attractor states currently encoded in the network (serving as a macro-scale prior) and the previous state, including external inputs (serving as likelihood in the Bayesian sense). The stochastic update, therefore, is equivalent to a Markov chain Monte Carlo (MCMC) sampling from this posterior distribution. The inverse temperature parameter β, in this regard, can be interpreted as the precision of the prior encoded in J. This is easy to conceptualize by considering the limiting case

of infinite precision, where the system simplifies to a binary-state Hopfield network ($\beta \to \infty$, $L(\beta\, u_i) \to \mathrm{sign}(u_i)$ on $[-1,1]$) that directly and deterministically converges to the (infinite-precision) prior, completely overriding the Bayesian likelihood (i.e., network input).

Free energy minimizing attractor networks as a model of large-scale brain dynamics

Taken together, the novel framework of free energy minimizing attractor networks not only motivates the use of a specific, emergent class of attractor networks as models for large-scale brain dynamics, but also provides a formal connection between these dynamics and Bayesian inference. The present study leverages this theoretical foundation. We aim to model large-scale brain dynamics as a free energy minimizing attractor network. According to our framework, such networks can be reconstructed from the activation time-series data measured in their nodes. Specifically, the weight matrix of the attractor network can be reconstructed as the negative inverse covariance matrix of the regional activation time series: $J = -\Lambda = -\Sigma^{-1}$, where Σ is the covariance matrix of the activation time series in all regions, and Λ is the precision matrix. For a detailed derivation, see Supplementary Information 4. Note that this approach can naturally be reduced to different "coarse-grainings" of the system, by pooling network nodes with similar functional properties. In the case of resting-state fMRI data, this corresponds to pooling network nodes into functional parcels. Drawing upon concepts such as the center manifold theorem (Wagner, 1989), it is posited that rapid, fine-grained dynamics at lower descriptive levels converge to lower-dimensional manifolds, upon which the system evolves via slower processes at coarser scales. It has been previously argued (Medrano *et al.*, 2024) that the temporal and spatial scales of fMRI data happen to align relatively well with the characteristic scales corresponding to meaningful large-scale "coarse-grainings" of brain dynamics.

Thus, we can reconstruct FEP-ANNs from functional connectivity data simply by considering the functional connectome (inverse covariance or partial correlation) as the coupling weights between the nodes of the network, which themselves correspond to brain regions (as defined by the chosen functional brain parcellation). We refer to such network models as functional connectivity-based attractor neural networks - **fcANN**s for short.

Having estimates of the weight matrix J of the attractor network, we can now rely on the deterministic and stochastic versions of the inference procedure (eq. (3)) in order to investigate this system. Running the deterministic update to a uniformly drawn sample of initial states, we can identify all attractor states of the network. The stochastic update, on the other hand, can be used to sample from the posterior distribution of the activity states, and thus serves as a generative computational model of the brain dynamics.

Testable predictions of the theoretical framework

Self-orthogonalization as a signature of free energy attractor networks

So far, we have only discussed free energy minimization in terms of the activity of the nodes of the network. However, free energy minimization also gives rise to a specific learning rule for the couplings J of the network. This learning rule is a specific local, incremental, contrastive (predictive coding-based) plasticity rule to adjust connection strengths:

$$\Delta J_{ij} \propto \underbrace{\sigma_i \sigma_j}_{\substack{\text{observed correlation} \\ \text{(Hebbian)}}} - \underbrace{L(b_i + \textstyle\sum_{k \neq i} J_{ik}\, \sigma_k)\, \sigma_j}_{\substack{\text{predicted correlation} \\ \text{(anti-Hebbian)}}} \tag{4}$$

A detailed derivation of the learning dynamics can be found in (Spisak & Friston, 2025) and Supplementary Information 2. In the present work, we do not implement this learning rule in our computational model, as the coupling weights J are reconstructed directly from the empirical fMRI activation time series data.

However, this specific learning rule has an important implication for the attractor states of the FEP-ANN: it will naturally drive them towards (approximate) orthogonality during learning. For a mathematical motivation of the mechanisms underlying this important property, termed self-orthogonalization, see (Spisak & Friston, 2025) and Supplementary Information 3. Self-orthogonalization is far from being a generic property of all attractor networks (and it is also not a consequence of the above formulated inference dynamics). It has, however, remarkable implications for the computational efficiency of the network and the robustness of its representations. Attractor networks with orthogonal attractor states, often termed the Kanter-Sompolinsky projector neural network (Kanter & Sompolinsky, 1987), are the computationally most efficient varieties of general attractor networks, with maximal memory capacity and perfect memory recall (without error). Importantly, in Kanter-Sompolinsky projector networks, **the eigenvectors of the coupling matrix and the attractors become equivalent**, providing an important signature for detecting such networks in empirical data.

Importantly, in the present study we reconstruct attractor networks from functional connectivity data (fcANNs) without relying on the learning rule of the FEP-ANN framework (eq. (4)) which imposes orthogonality on the attractors. Thus, if, in these empirically reconstructed fcANNs, an alignment between the eigenvectors of the coupling matrix and the attractors is observed, it can be considered strong evidence that the system approximates a Kanter–Sompolinsky projector network. As FEP-ANNs - together with some other, related models (e.g., "dreaming neural networks" (Hopfield *et al.*, 1983; Plakhov, n.d.; Dotsenko & Tirozzi, 1991; Fachechi *et al.*, 2019)) - provide a plausible and mathematically rigorous mechanistic model for the emergence of architectures approximating Kanter–Sompolinsky projector networks through biologically plausible local learning rules, this alignment between the eigenvectors of the coupling matrix and the attractors can be considered a signature of an underlying FEP-ANN. We will directly test this prediction in the present study, by investigating the orthogonality of the attractor states of the fcANN model reconstructed from empirical fMRI data.

Convergence, multistability, biological plausibility and prediction capacity

Beyond (approximate) attractor orthogonality, our framework provides additional testable predictions. If the functional connectome can indeed be considered a proxy for the coupling weights J of an underlying attractor network, we can expect that (i) the reconstructed fcANN model will exhibit multiple stable attractor states, with large basins and biologically plausible spatial patterns, (ii) the relaxation dynamics of the reconstructed model will display fast convergence to attractor states, and (iii) the stochastic relaxation dynamics yield an efficient generative model of the empirical resting-state brain dynamics as well as perturbations thereof caused either by external inputs (stimulations and tasks) or pathologies.

Research questions

We have structured the present work around 7 research questions we address in the present study:

Q1 - Is the brain an approximate K-S projector ANN (FEP-ANN prediction)?
We test whether fcANN-derived brain attractor states closely resemble the eigenvectors of the functional connectome matrix, in contrast to null models based on temporally phase-randomized time series data (preserving the frequency spectrum and the temporal autocorrelation of data, but destroying conditional dependencies across regions), denoted as **NM1**. Furthermore, in a supplementary analysis, we quantify the similarity of the functional connectome to the weights of an optimal Kanter–Sompolinsky (K-S) network with the same eigenvectors. The similarity (cosine similarity) is contrasted against repeating the same approach on permuted coupling matrices (but retaining symmetry, **NM2**)
Q2 - Is the functional connectome well suited to function as an attractor network?
We contrast the convergence properties of fcANN deterministic relaxation dynamics with null models with permuted coupling weights (preserving symmetry, sparsity and weight distributions, destroying topological structure) *NM2*.
Q3 - What are the optimal parameters for the fcANN model?
The number of attractor states is a function of the inverse temperature parameter β. For simplicity, we fix $\beta = 0.04$ (4 attractor states) in the current analysis. We perform a rough optimization of the noise parameter ϵ by benchmarking the fcANN's ability to capture non-Gaussian conditional distributions in the data. This is benchmarked by computing a Wasserstein distance between the distributions of empirical and simulated data and contrasting it to the null model of a multivariate normal distribution with covariance matched to that of the empirical data (**NM3**, representing the case of Gaussian-only conditionals).
Q4 - Do fcANNs display biologically plausible attractor states?
We qualitatively demonstrate that attractor states obtained with different inverse temperature parameters β and different noise levels (ϵ) exhibit large basins and that these attractor states exhibit spatial patterns consistent with known large-scale brain systems.
Q5 - Can fcANNs reproduce the characteristics of resting-state brain activity?
We compare how well fcANN attractor states explain variance in unseen (in- and out-of-sample) empirical time series data, relative to the principal components of the empirical data itself. Statistical significance is evaluated via bootstrapping. Furthermore, we compare various characteristics (state occupancy, distribution, temporal trajectory) of the data generated by fcANNs via stochastic updates to empirical resting-state data. As null models, we use covariance-matched multivariate normal distributions (NM3)
Q6 - Can resting-state fMRI-based fcANNs predict large-scale brain dynamics elicited by tasks or stimuli?
We test whether fcANNs initialized from resting-state functional connectomes and perturbed with weak, condition-specific control signals predict task-evoked large-scale dynamics (pain vs. rest; up- vs. down-regulation). We compare simulated and empirical differences on the fcANN projection and flow fields during stochastic updates. As a null model, we use condition-label shuffling (NM5).
Q7 - Can resting-state fMRI-based fcANNs predict altered brain dynamics in clinical populations?

We test whether fcANNs initialized with group-level resting-state connectomes from autism spectrum disorder (ASD) patients and typically developing controls (TDC) predict observed group differences in dynamics (state occupancy, attractor-basin activations, flow fields). We compare fcANN-generated dynamics between ASD- and TDC-initialized models and evaluate similarity to empirical contrasts. As a null model, we use group-label shuffling (NM5).
For a summary of null modelling approaches and research questions, see Table 1 and Table 2.

Table 1: Null models applied in the present study

Short name	***Brief description***	***Invariant to***	***Destroys***
NM1: *Temporal phase randomization*	Phase-randomize time series data independently for each region; recalculate connectivity.	Time-series power spectrum and autocorrelation	Conditional dependencies across regions
NM2: *Symmetry-preserving matrix permutation*	Shuffle off-diagonal entries of $\mathbf{J}$ while keeping symmetry	Weight distribution and symmetry	Topological structure, clusteredness
NM3: *Covariance-matched Gaussian*	Draw time frames from a multivariate normal with covariance equal to the functional connectome's covariance	Gaussian conditionals	Nonlinear and non-Gaussian conditionals, temporal autocorrelation
NM4: *Temporal order permutation*	Randomly permute time-frame order within runs; used for flow analyses	Spatial autocorrelation	Temporal autocorrelation
NM5 *Condition shuffling*	permute condition labels, either within participant (e.g., pain vs. rest; up- vs. down-regulation) or between participant (shuffle patient vs. control labels)	Marginal distributions and overall data structure	Condition-specific associations and effects

Table 2: Research questions, methodological approaches, and the corresponding null models

Research Question	***Methodological Approach***	***Null Model***
Q1. *Is the brain an approximate K-S projector ANN (FEP-ANN prediction)?*	Compare eigenvectors of coupling matrix with attractor states	NM1-2
Q2. *Is the functional connectome well suited to function as an attractor network?*	Measure iterations to convergence in deterministic relaxation	NM2
Q3. *What are the optimal parameters for the fcANN model?*	We fix $\beta = 0.04$ (4 attractor states) for simplicity. We perform a rough optimization of the noise parameter ϵ in stochastic relaxation to match empirical data distribution.	NM3
Q4. *Do fcANNs display biological plausible attractor states?*	Identify attractor states, report basin sizes and assess spatial patterns with different inverse temperature parameters and noise levels	qualitative
Q5. *Can fcANNs reproduce the characteristics of resting-state brain activity?*	Compare stochastic dynamics (state occupancy, distribution, temporal trajectory) with empirical resting state data	NM3-4
Q6. *Can resting-state fMRI-based fcANNs predict large-scale brain dynamics elicited by tasks or stimuli?*	Contrast pain vs. rest dynamics with data generated by fcANNs and pain-associated control signal	NM5
Q7. *Can resting-state fMRI-based fcANNs predict altered brain dynamics in clinical populations?*	Contrast autism spectrum disorder patients vs. typically developing control participants' observed brain dynamics with data generated by fcANNs initialized with the respective functional connectomes	NM5

Results

Functional connectivity-based attractor networks (fcANNs) as a model of brain dynamics

First, we constructed a functional connectivity-based attractor network (fcANN) based on resting-state fMRI data in a sample of n=41 healthy young participants (study 1). Details are described in the Methods. In brief, we estimated interregional activity flow (Cole *et al.*, 2016; Ito *et al.*, 2017) as the study-level average of regularized partial correlations among the resting-state fMRI time series of m=122 functional parcels of the BASC brain atlas (see Methods for details). We then used the standardized functional connectome as the J_{ij} weights of a fully connected recurrent fcANN model, see Methods).

Next, we applied the deterministic relaxation procedure to a large number of random initializations (n=100000) to obtain all possible attractor states of the fcANN in study 1 (Figure 2A). Consistent with theoretical expectations, we observed that increasing the inverse temperature parameter β led to an increasing number of attractor states (Figure 2E, left, Supplementary Figure 5), appearing in symmetric pairs (i.e. $\sigma_i^{(1)} = -\sigma_i^{(2)}$, see Figure 2G).

To test research question **Q1**, we matched the eigenvectors of the coupling matrix to the attractor state with which they exhibit the highest correlation. We compared eigenvector–attractor correlations with a null model based on phase-randomized surrogate time-series data (NM1). We found that the eigenvectors of the coupling matrix and the attractor states are significantly more strongly aligned (as measured with Pearson's correlation coefficient) than those in the null model (two-sided empirical permutation test; 1,000 permutations; correlations and p-values for the first six eigenvector–attractor pairs are reported in Figure 2A), providing evidence that large-scale brain organization approximates a Kanter–Sompolinsky projector network architecture (Figure 2A). Eigenvectors with the highest eigenvalues tended to be aligned with the attractor states with the highest fractional occupancy (ratio of time spent on their basins during simulations with stochastic relaxation; see Figure 2F). No such pattern was observed in the null model (Figure 2B). Further evidence for the functional connectome's close resemblance to a Kanter–Sompolinsky projector network is provided by the orthogonality of the attractor states to each other (Figure 2C) and additional analyses reported in Supplementary Figure 1.

Next, to support the visualization of further analyses, we constructed a simplified, 2-dimensional visual representation of fcANN dynamics, which we apply throughout the remaining manuscript as a high-level visual summary. This 2-dimensional visualization, referred to as the *fcANN* projection, is based on the first two principal components (PCs) of the states sampled from the stochastic relaxation procedure (Figure 2D–F and Supplementary Figure 2). On this simplified visualization, we observed a clear separation of the attractor states (Figure 2D), with the two symmetric pairs of attractor states located at the extremes of the first and second PC. To map the attractors' basins on the space spanned by the first two PCs (Figure 2C), we obtained the attractor state of each point visited during the stochastic relaxation and fit a multinomial logistic regression model to predict the attractor state from the first two PCs. The resulting model accurately predicted attractor states of arbitrary brain activity patterns, achieving a cross-validated accuracy of 96.5% (two-sided empirical permutation $p<0.001$; 1,000 label permutations within folds). This allows us to visualize attractor basins on this 2-dimensional projection by delineating the decision boundaries obtained from this model (Supplementary Figure 2 as black lines in Figure 2G-H). In the rest of the manuscript, we use this 2-dimensional fcANN projection depicted on (Figure 2D-H) as a simplified visual representation of brain dynamics.

Panel D on Figure 2 uses the fcANN projection to visualize the conventional Hopfield relaxation procedure. It depicts the trajectory of individual activation maps (sampled randomly from the time series data in study 1) until converging to one of the four attractor states. Panel E shows that the system does not converge to an attractor state with the stochastic relaxation procedure. The resulting path is still influenced by the attractor states' "gravitational pull", resulting in multistable dynamics that resemble the empirical time series data (example data on panel F).

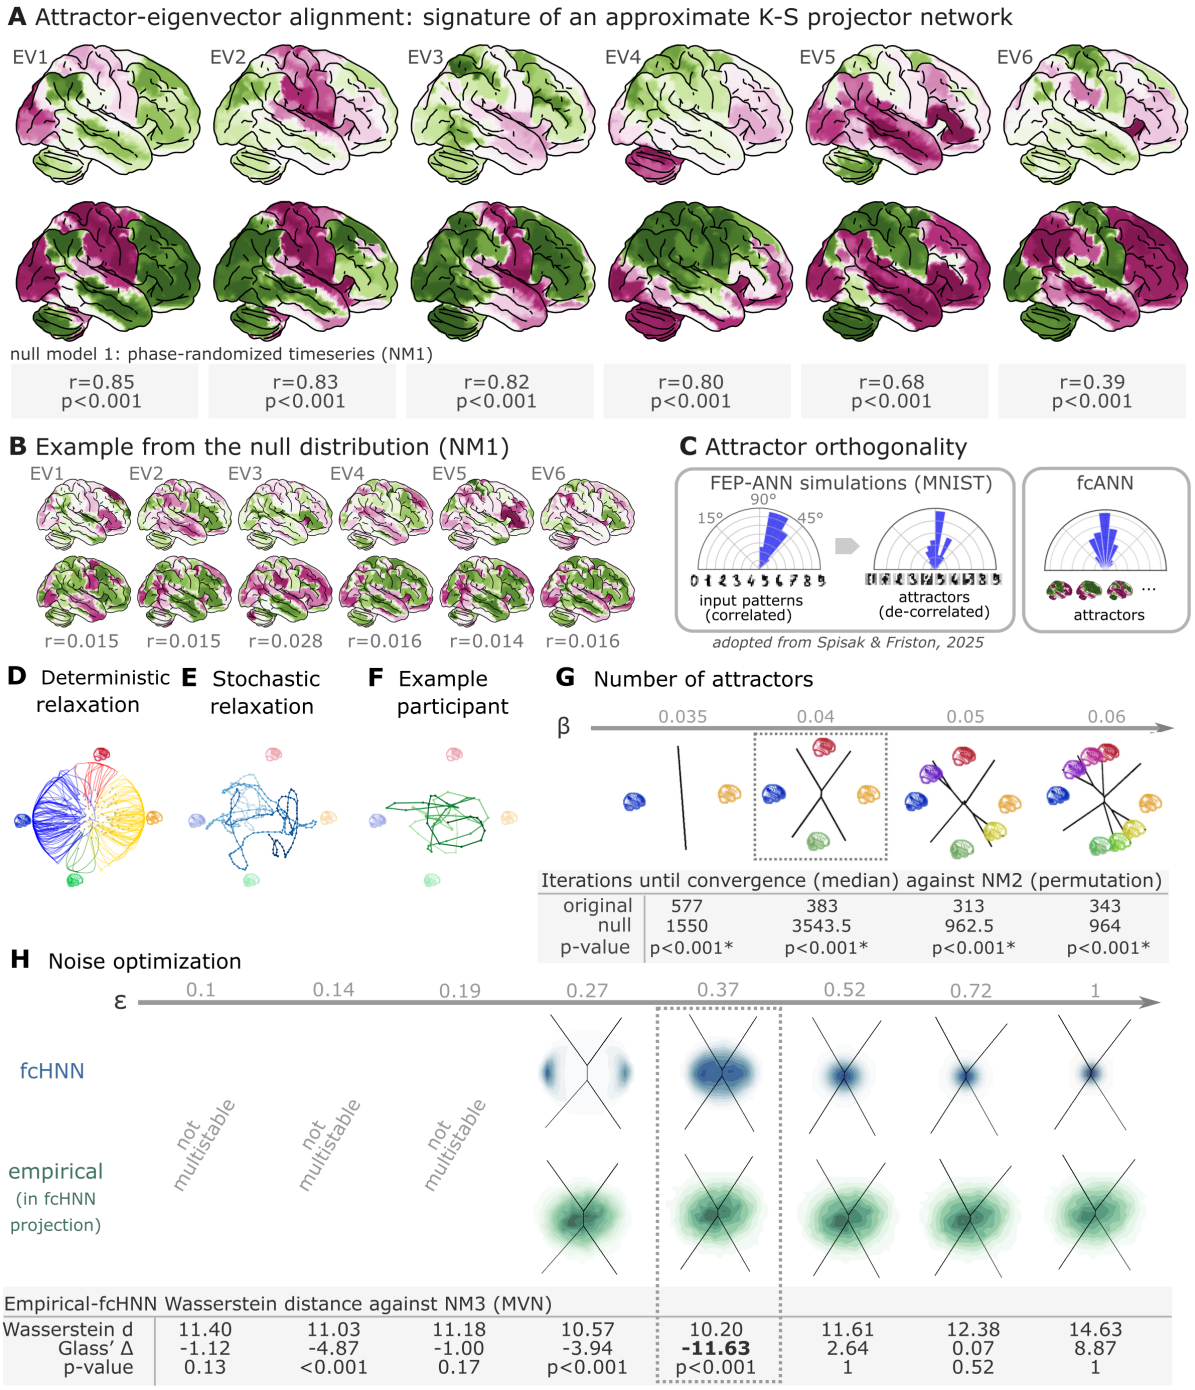

Figure 2: Attractor states and state-space dynamics of connectome-based Hopfield networks **A** Numbers under each pair report Pearson correlation and two-sided p-values based on 1,000 surrogate data realizations, generated by phase-randomizing the true time series and recomputing the connectivity matrix. For the comprehensive results of the eigenvector–attractor alignment analysis (including a supplementary analysis on weight similarity to the analogous Kanter–Sompolinsky projector network) see Supplementary Figure 1. **B** Example matches from a single permutation of the permutation-based null distribution. For each symmetry-preserving permutation of J, we recomputed the corresponding eigenvectors and attractors and re-matched them. The maps are visibly mismatched and correlations are near zero, illustrating the null against which the empirical correlations in panel A are evaluated. **C** Left panel: Free-energy-minimizing attractor networks have been shown to establish approximately orthogonal attractor states (right), even when presented with correlated patterns (left, adapted from Spisak & Friston (2025)). fcANN analysis reveals that the brain also exhibits approximately orthogonal attractors. On all three polar plots, pairwise angles between attractor states are shown. Angles concentrating around 90° in the empirical fcANN are consistent with predictions of free-energy-minimizing (Kanter–Sompolinsky-like) networks. (Note, however, that in high-dimensional spaces, random vectors would also tend to be approximately orthogonal.) **D** The fcANN of study 1 seeded with real activation maps (gray dots) of an example participant. All activation

maps converge to one of the four attractor states during the deterministic relaxation procedure (without noise) and the system reaches equilibrium. Trajectories are colored by attractor state. **E** Illustration of the stochastic relaxation procedure in the same fcANN model, seeded from a single starting point (activation pattern). With stochastic relaxation, the system no longer converges to an attractor state, but instead traverses the state space in a way restricted by the topology of the connectome and the "gravitational pull" of the attractor states. The shade of the trajectory changes with increasing number of iterations. The trajectory is smoothed with a moving average over 10 iterations for visualization purposes. **F** Real resting state fMRI data of an example participant from study 1, plotted on the fcANN projection. The shade of the trajectory changes with an increasing number of iterations. The trajectory is smoothed with a moving average over 10 iterations for visualization purposes. **G** Consistent with theoretical expectations, we observed that increasing the inverse temperature parameter β led to an increasing number of attractor states, emerging in a nested fashion (i.e. the basin of a new attractor state is fully contained within the basin of a previous one). When contrasting the functional connectome-based ANN with a null model based on symmetry-retaining permuted variations of the connectome (NM2), we found that the topology of the original (unpermuted) functional brain connectome makes it significantly better suited to function as an attractor network than the permuted null model. Table contains the median number of iterations until convergence for the original and permuted connectomes for different temperature parameters β and the p-value derived from a one-sided Wilcoxon signed-rank test (i.e. a non-parametric paired test) comparing the iteration values for each random null instance (1,000 pairs) to the iteration number observed with the original matrix and the same random input; with the null hypothesis that the empirical connectome converges in fewer iterations than the permuted connectome. **H** We optimized the noise parameter ϵ of the stochastic relaxation procedure for 8 different ϵ values over a logarithmic range between $\epsilon = 0.1$ and 1 and contrasted the similarity (Wasserstein distance) between the 122-dimensional distribution of the empirical and the fcANN-generated data against null data generated from a covariance-matched multivariate normal distribution (1000 surrogates). We found that the fcANN reached multistability with $\epsilon > 0.19$ and provided the most accurate reconstruction of the real data with $\epsilon = 0.37$, as compared with its accuracy in retaining the null data, suggesting that the fcANN model is capable of capturing non-Gaussian conditionals in the data. Glass's Delta quantifies the distance from the null mean, expressed in units of null standard deviation.

In study 1, we investigated the convergence process of the fcANN (research question **Q2**) and contrasted it with a null model based on permuted variations of the connectome (while retaining the symmetry of the matrix, NM2). This null model preserves the sparseness and the degree distribution of the connectome, but destroys its topological structure (e.g. clusteredness). We found that the topology of the original (unpermuted) functional brain connectome makes it significantly better suited to function as an attractor network than the permuted null model. For instance, with $\beta = 0.04$, the median iteration number for the original and permuted fcANNs to reach convergence was 383 and 3543.5 iterations, respectively (Figure 2G, Supplementary Figure 6). Similar results were observed, independent of the inverse temperature parameter β. We set the temperature parameter for the rest of the paper to a value of $\beta = 0.04$, resulting in 4 distinct attractor states. The primary motivation for selecting $\beta = 0.04$ was to reduce the computational burden and the interpretational complexity for further analyses. However, as with increasing temperature attractor states emerge in a nested fashion, we expect that the results of the following analyses would be, although more detailed, qualitatively similar with higher β values.

Next, in line with research question **Q3**, we optimized the noise parameter ϵ of the stochastic relaxation procedure for 8 different ϵ values over a logarithmic range between $\epsilon = 0.1$ and 1 and contrasted the similarity (Wasserstein distance) between the 122-dimensional distribution of the empirical and the fcANN-generated data against null data generated from a covariance-matched multivariate normal distribution (1000 surrogates). We found that the fcANN reached multistability with $\epsilon > 0.19$ and provided the most accurate reconstruction of the non-Gaussian conditional dependencies in the real data with $\epsilon = 0.37$, as compared to its accuracy in retaining the covariance-matched multivariate Gaussian null data (NM3 Figure 2H; Wasserstein distance: 10.2, Glass's Delta (distance from null mean, expressed in units of null standard deviation): -11.63, $p<0.001$ one-sided). Based on this coarse optimization procedure, we set $\epsilon = 0.37$ for all subsequent analyses.

Reconstruction of resting state brain dynamics

Next, we visualized and qualitatively assessed the neuroscientific relevance of the spatial patterns of the obtained attractor states (**Q4**, Figure 3A), and found that they closely resemble previously described large-scale brain systems. The spatial patterns associated with first pair of attractors (mapped on PC1 on the 2-dimensional projection, horizontal axis, e.g. on Figure 2D-H) show a close correspondence to two commonly described complementary brain systems, that have been previously found in anatomical, functional, developmental, and evolutionary hierarchies, as well as gene expression, metabolism, and blood flow, (see Sydnor *et al.* (2021) for a review), and reported under various names, like intrinsic and extrinsic systems (Golland *et al.*, 2008), Visual-Sensorimotor-Auditory and Parieto-Temporo-Frontal "rings" (Cioli *et al.*, 2014), "primary" brain substates (Chen *et al.*, 2018), unimodal-to-transmodal principal gradient (Margulies *et al.*, 2016; Huntenburg *et al.*, 2018) or sensorimotor-association axis (Sydnor *et al.*, 2021). A common interpretation of these two patterns is that they represent (i) an "intrinsic" system for higher-level internal context, commonly referred to as the *default mode*

network (Raichle *et al.*, 2001) and (ii) an anti-correlated "extrinsic" system linked to the immediate sensory environment, showing similarities to the recently described "action mode network" (Dosenbach *et al.*, 2025). The other pair of attractors - spanning an approximately orthogonal axis - resemble patterns commonly associated with perception–action cycles (Fuster, 2004), and described as a gradient across sensorimotor modalities (Huntenburg *et al.*, 2018), recruiting regions associated with active (e.g. motor cortices) and perceptual inference (e.g., visual areas).

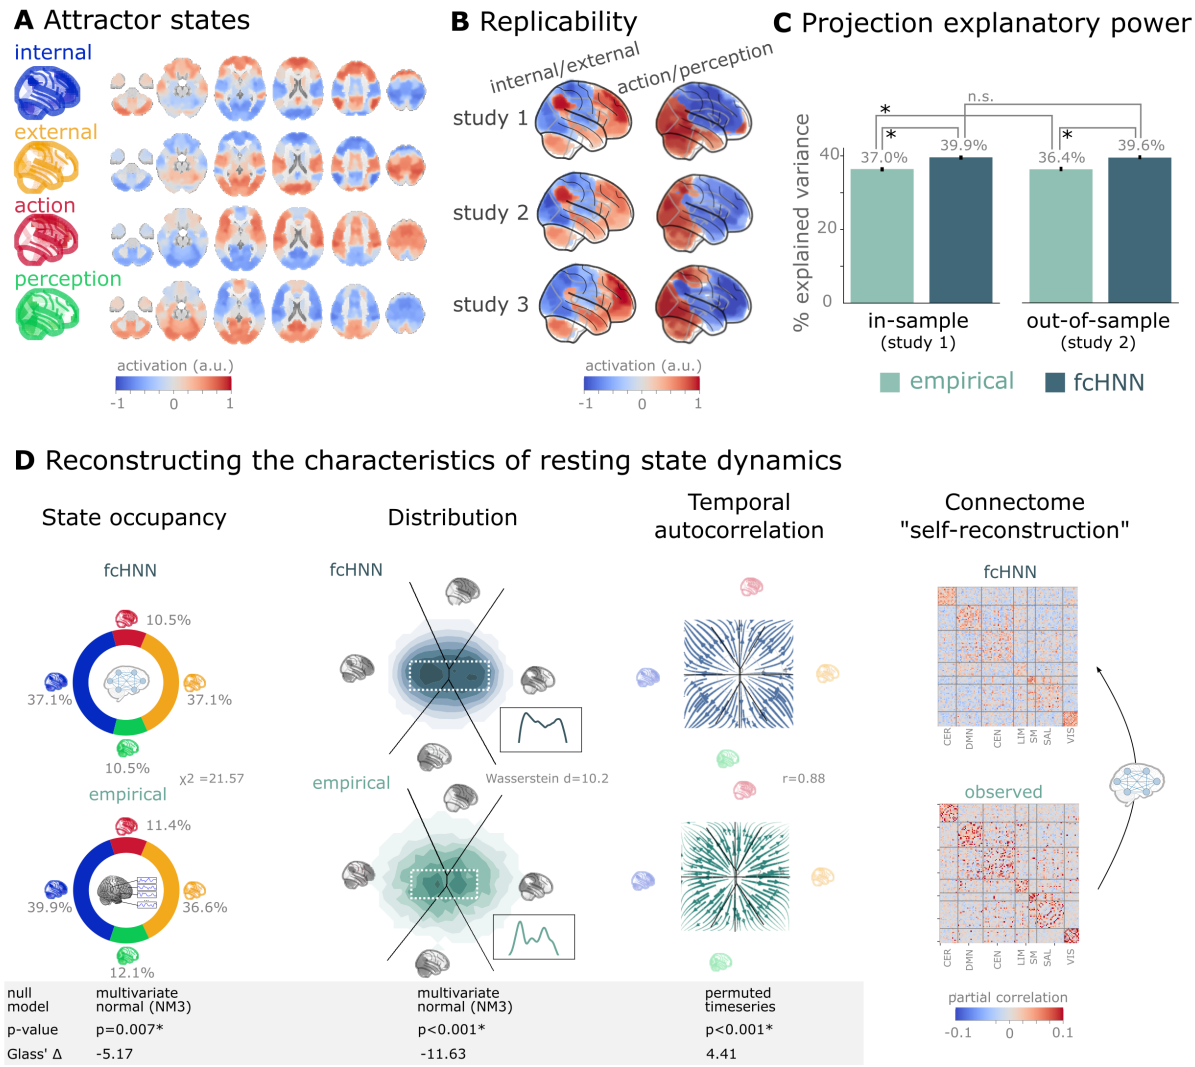

Figure 3: Connectome-based attractor networks reconstruct characteristics of real resting-state brain activity.

A The four attractor states of the fcANN model from study 1 reflect brain activation patterns with high neuroscientific relevance, representing sub-systems previously associated with "internal context" (blue), "external context" (yellow), "action" (red) and "perception" (green) (Golland *et al.*, 2008; Cioli *et al.*, 2014; Chen *et al.*, 2018; Fuster, 2004; Margulies *et al.*, 2016; Dosenbach *et al.*, 2025). **B** The attractor states show excellent replicability in two external datasets (study 2 and 3, overall mean correlation 0.93). **C** The first two PCs of the fcANN state space (the "fcANN projection") explain significantly more variance (two-sided percentile bootstrap $p<0.0001$ on ΔR^2, 100 resamples) in the real resting-state fMRI data than principal components derived from the real resting-state data itself and generalizes better (two-sided percentile bootstrap $p<0.0001$) to out-of-sample data (study 2). Error bars denote 99% percentile bootstrapped confidence intervals (100 resamples). **D** The fcANN analysis reliably predicts various characteristics of real resting-state fMRI data, such as the fraction of time spent on the basis of the four attractors (first column, $p=0.007$, contrasted to the multivariate normal null model NM3), the distribution of the data on the fcANN-projection (second column, $p<0.001$, contrasted to the multivariate normal null model NM3) and the temporal autocorrelation structure of the real data (third column, $p<0.001$, contrasted to a null model based on permuting time-frames). The latter analysis was based on flow maps of the mean trajectories (i.e. the characteristic timeframe-to-timeframe transition direction) in fcANN-generated data, as compared to a shuffled null model representing zero temporal autocorrelation. For more details, see Methods. Furthermore, we demonstrate that - in line with the theoretical expectations - fcANNs "leak" their weights during stochastic inference (rightmost column): the time series

resulting from the stochastic relaxation procedure mirror the covariance structure of the functional connectome the fcANN model was initialized with. While the "self-reconstruction" property in itself does not strengthen the face validity of the approach (no unknown information is reconstructed), it is a strong indicator of the model's construct validity; i.e. that systems that behave like the proposed model inevitably "leak" their weights into the activity time series.

The discovered attractor states demonstrate high replicability across the discovery dataset (study 1) and two independent replication datasets (study 2 and 3, Figure 3C; overall mean Pearson's correlation 0.93, pooled across datasets and attractor states). In a supplementary analysis, we have also demonstrated the robustness of fcANNs to imperfect functional connectivity measures: fcANNs were found to be significantly more robust to noise added to the coupling matrix than nodal strength scores (used as a reference with the same dimensionality; see Supplementary Figure 10 for details).

Further analysis in study 1 showed that connectome-based attractor models accurately reconstructed multiple characteristics of true resting-state data (**Q5**). First, the two "axes" of the fcANN projection (corresponding to the first four attractors) accounted for a substantial amount of variance in the real resting-state fMRI data in study 1 (mean $R^2 = 0.399$) and generalized well to out-of-sample data (study 2, mean $R^2 = 0.396$) (Figure 3E). The variance explained by the attractors significantly exceeded that of the first two PCs derived directly from the real resting-state fMRI data itself ($R^2 = 0.37$ and 0.364 for in- and out-of-sample analyses). PCA - by identifying variance-heavy orthogonal directions - aims to explain the highest amount of variance possible in the data (with the assumption of Gaussian conditionals). While empirical attractors are closely aligned to the PCs (i.e. eigenvectors of the inverse covariance matrix), the alignment is only approximate. Here we quantified whether attractor states are a better fit to the unseen data than the PCs. Obviously, due to the otherwise strong PC–attractor correspondence, this is expected to be only a small improvement. However, this provides important evidence for the validity of our framework, as - together with our analysis addressing Q3 - it shows that attractors are not just a complementary, perhaps "noisier" variety of the PCs, but a "substrate" that generalizes better to unseen data than the PCs themselves.

Second, during stochastic relaxation, the fcANN model was found to spend approximately three-quarters of the time on the basins of the first two attractor states and one-quarter on the basins of the second pair of attractor states (approximately equally distributed between pairs). We observed similar temporal occupancies in the real data (Figure 3D left column), statistically significant against a covariance-matched multivariate Gaussian null model (NM3, 1,000 surrogates each; observed $\chi^2 = 21.57$, $p<0.001$; Glass $\Delta = -5.17$; see Supplementary Figure 7 for details and for an alternative null model based on spatial phase-randomization). Fine-grained details of the distribution with bimodal appearance, observed in the real resting-state fMRI data were also convincingly reproduced by the fcANN model (Figure 3F and Figure 2D, second column).

Third, not only spatial activity patterns but also time series generated by the fcANN are similar to empirical time series data. Next to the visual similarity shown on Figure 2E and F, we observed a statistically significant similarity between the average trajectories of fcANN-generated and real time series "flow" (i.e. the characteristic timeframe-to-timeframe transition direction, Pearson's $r = 0.88$, $p<0.001$, Glass $\Delta = 4.41$), as compared to null models of zero temporal autocorrelation (randomized timeframe order; two-sided empirical permutation test on Pearson's r with 1,000 permutations; Figure 3D, third column; Methods).

Finally, we have demonstrated that - as expected from theory - fcANNs generate signal that preserves the covariance structure of the functional connectome they were initialized with, indicating that dynamic systems of this type (including the brain) inevitably "leak" their underlying structure into the activity time series, strengthening the construct validity of our approach (Figure 3D).

An explanatory framework for task-based brain activity

Next to reproducing various characteristics of spontaneous brain dynamics, fcANNs can also be used to model responses to various perturbations (**Q6**). We obtained task-based fMRI data from a study by Woo *et al.* (2015) (study 4, n=33, see Figure 3), investigating the neural correlates of pain and its self-regulation.

We found that activity changes due to pain (taking into account hemodynamics, see Methods) were characterized on the fcANN projection by a shift toward the attractor state of action/execution (NM5: two-sided permutation test on the L2 norm of the mean projection difference; 1,000 within-participant label swaps; $p<0.001$; Glass's $\Delta = 4.34$; Figure 4A, left). Energies, as defined by the fcANN, were also significantly different between the two conditions (NM5: two-sided permutation test on absolute energy difference; 1,000 label swaps; $p<0.001$; Glass's $\Delta = 3.14$), with higher energies during pain stimulation.

When participants were instructed to up- or downregulate their pain sensation (resulting in increased and decreased pain reports and differential brain activity in the nucleus accumbens, NAc (see Woo *et al.* (2015) for details), we observed further changes in the location of momentary brain activity patterns on the fcANN projection

(two-sided permutation test on the L2 norm of the mean projection difference; 1,000 label swaps; p<0.001; Glass's Δ = 4.1; Figure 4A, right), with downregulation pulling brain dynamics toward the attractor state of internal context and perception. Interestingly, self-regulation did not trigger significant energy changes (two-sided permutation test on absolute energy difference; 1,000 label swaps; p=0.37; Glass's Δ = 0.4).

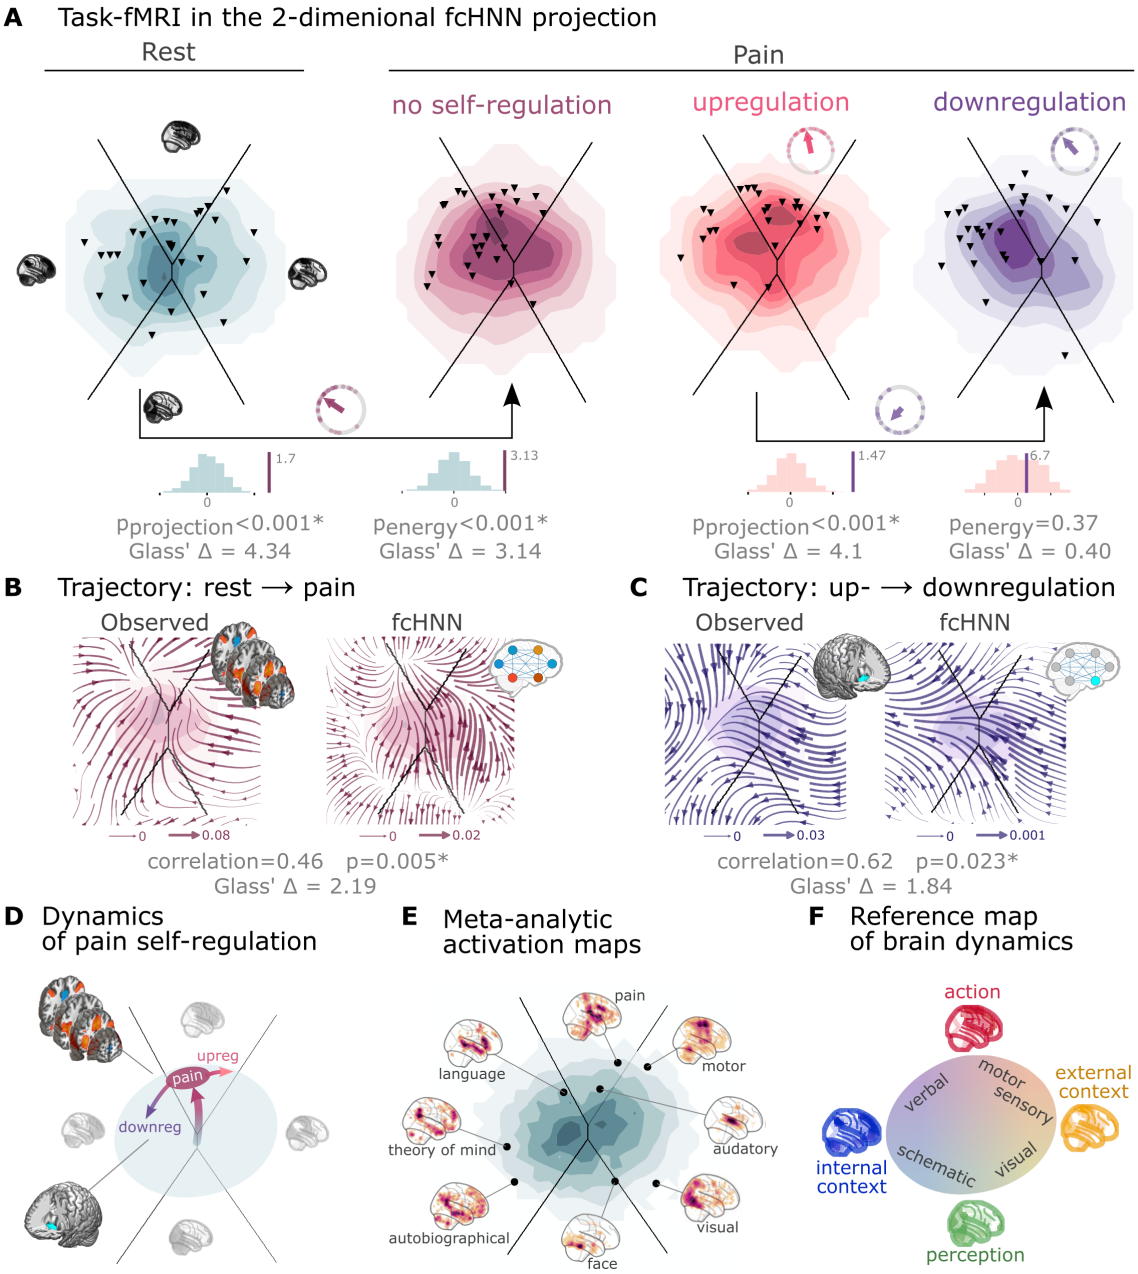

Figure 4: Functional connectivity-based attractor networks reconstruct real task-based brain activity. **A** Functional MRI time-frames during pain stimulation from study 4 (second fcANN projection plot) and self-regulation (third and fourth) are distributed differently on the fcANN projection than brain substates during rest (first projection, permutation test, p<0.001 for all). Energies, as defined by the Hopfield model, are also significantly different between rest and the pain conditions (permutation test, p<0.001), with higher energies during pain stimulation. Triangles denote participant-level mean activations in the various blocks (corrected for hemodynamics). Small circle plots show the directions of the change for each individual (points) as well as the mean direction across participants (arrow), as compared to the reference state (downregulation for the last circle plot, rest for all other circle plots). **B** Flow-analysis (difference in the average timeframe-to-timeframe transition direction) reveals a nonlinear difference in brain dynamics during pain and rest (left). When introducing weak pain-related signal in the fcANN model during stochastic relaxation, it accurately reproduces these nonlinear flow differences (right). **C** Simulating activity in the Nucleus Accumbens (NAc) (the region showing significant activity differences in Woo *et al.* (2015)) reconstructs the observed nonlinear flow difference between up- and downregulation (left). **D** Schematic representation of brain dynamics during pain and its up- and

downregulation, visualized on the fcANN projection. In the proposed framework, pain does not simply elicit a direct response in certain regions, but instead, shifts spontaneous brain dynamics towards the "action" attractor, converging to a characteristic "ghost attractor" of pain. Down-regulation by NAc activation exerts force towards the attractor of internal context, leading to the brain less frequent "visiting" pain-associated states. **E** Visualizing meta-analytic activation maps (see Supplementary Table 2 for details) on the fcANN projection captures intimate relations between the corresponding tasks and **F** serves as a basis for a fcANN-based theoretical interpretative framework for spontaneous and task-based brain dynamics. In the proposed framework, task-based activity is not a mere response to external stimuli in certain brain locations but a perturbation of the brain's characteristic dynamic trajectories, constrained by the underlying functional connectivity. From this perspective, "activity maps" from conventional task-based fMRI analyses capture time-averaged differences in these whole brain dynamics.

Next, we conducted a "flow analysis" on the fcANN projection, quantifying how the average timeframe-to-timeframe transition direction differs on the fcANN projection between conditions (see Methods). This analysis unveiled that during pain (Figure 4B, left side), brain activity tends to gravitate toward a distinct point on the projection on the boundary of the basins of the internal and action attractors, which we term the "ghost attractor" of pain (similar to Vohryzek *et al.* (2020)). In case of downregulation (as compared to upregulation), brain activity is pulled away from the pain-related "ghost attractor" (Figure 4C, left side), toward the attractor of internal context.

Our fcANN was able to accurately reconstruct these nonlinear dynamics by adding a small amount of realistic "control signal" (similarly to network control theory, see e.g. Liu *et al.* (2011) and Gu *et al.* (2015)). To simulate the alterations in brain dynamics during pain stimulation, we acquired a meta-analytic pain activation map (Zunhammer *et al.*, 2021) (n=603) and incorporated it as a control signal added to each iteration of the stochastic relaxation procedure. The ghost attractor found in the empirical data was present across a relatively wide range of signal-to-noise (SNR) values (Supplementary Figure 8). Results with SNR=0.005 are presented in Figure 4B, right side (Pearson's r = 0.46; two-sided permutation p=0.005 based on NM5: randomizing conditions on a per-participant basis; 1,000 permutations; Glass's $\Delta = 2.19$).

The same model was also able to reconstruct the observed nonlinear differences in brain dynamics between the up- and downregulation conditions (Pearson's r = 0.62; p=0.023 based on two-sided permutation test NM5: randomly shuffling conditions in a per-participant basis; 1,000 permutations; Glass's $\Delta = 1.84$) without any further optimization (SNR=0.005, Figure 4C, right side). The only change we made to the model was the addition (downregulation) or subtraction (upregulation) of control signal in the NAc (the region in which (Woo *et al.*, 2015) observed significant changes between up- and downregulation), introducing a signal difference of ΔSNR=0.005 (the same value we found optimal in the pain-analysis). Results were reproducible with lower NAc SNRs, too (Supplementary Figure 9).

To provide a comprehensive picture on how tasks and stimuli other than pain map onto the fcANN projection, we obtained various task-based meta-analytic activation maps from Neurosynth (see Methods) and plotted them on the fcANN projection (Figure 4E). This analysis reinforced and extended our interpretation of the four investigated attractor states and shed more light on how various functions are mapped on the axes of internal vs. external context and perception vs. action. In the coordinate system of the fcANN projection, visual processing is labeled "external-perception", sensory-motor processes fall into the "external-active" domain, language, verbal cognition and working memory belongs to the "internal-active" region and long-term memory as well as social and autobiographic schemata fall into the "internal-perception" regime (Figure 4F).

Clinical relevance

To demonstrate fcANN models' potential to capture altered brain dynamics in clinical populations (**Q7**), we obtained data from n=172 autism spectrum disorder (ASD) and typically developing control (TDC) individuals, acquired at the New York University Langone Medical Center, New York, NY, USA (NYU) and generously shared in the Autism Brain Imaging Data Exchange dataset (study 7: ABIDE, (Di Martino *et al.*, 2014)). After excluding high-motion cases (with the same approach as in studies 1–4, see Methods), we visualized the distribution of time-frames on the fcANN-projection separately for the ASD and TDC groups (Figure 5A). First, we assigned all timeframes to one of the 4 attractor states with the fcANN from study 1 and found several significant differences in the mean activity on the attractor basins (see Methods) of the ASD group as compared to the respective controls (Figure 5B). Strongest differences were found on the "action-perception" axis (Table 3), with increased activity of the sensory-motor and middle cingular cortices during "action-execution" related states and increased visual and decreased sensory and auditory activity during "perception" states, likely reflecting the widely acknowledged, yet poorly understood, perceptual atypicalities in ASD (Hadad & Schwartz, 2019). ASD related changes in the internal-external axis were characterized by more involvement of the posterior cingulate, the precuneus, the nucleus accumbens, the dorsolateral prefrontal cortex (dlPFC), the cerebellum (Crus II, lobule VII) and inferior temporal regions during activity of the internalizing subsystem

(Table 3). While similar, default mode network (DMN)-related changes have often been attributed to an atypical integration of information about the "self" and the "other" (Padmanabhan *et al.*, 2017), a more detailed fcANN-analysis may help to further disentangle the specific nature of these changes.

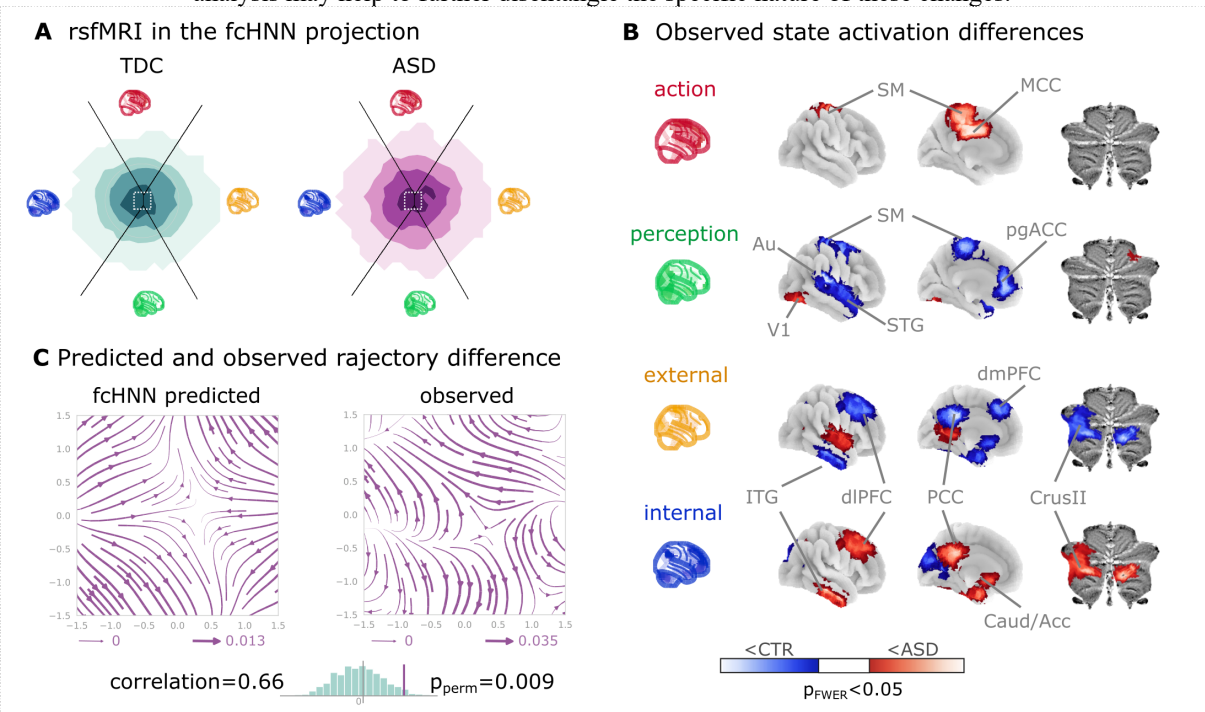

Figure 5: Connectome-based Hopfield analysis of autism spectrum disorder.
A The distribution of time-frames on the fcANN-projection separately for ASD patients and typically developing control (TDC) participants.
B We quantified attractor state activations in the Autism Brain Imaging Data Exchange datasets (study 7) as the individual-level mean activation of all time-frames belonging to the same attractor state. This analysis captured alterations similar to those previously associated with ASD-related perceptual atypicalities (visual, auditory and somatosensory cortices) as well as atypical integration of information about the "self" and the "other" (default mode network regions). All results are corrected for multiple comparisons across brain regions and attractor states (122×4 comparisons) with Bonferroni correction. See Table 3 and Supplementary Figure 11 for detailed results.
C The comparison of data generated by fcANNs initialized with ASD and TDC connectomes, respectively, revealed a characteristic pattern of differences in the system's dynamics, with increased pull towards (and potentially a higher separation between) the action and perception attractors and a lower tendency of trajectories going towards the internal and external attractors.
Abbreviations: MCC: middle cingulate cortex, ACC: anterior cingulate cortex, pg: perigenual, PFC: prefrontal cortex, dm: dorsomedial, dl: dorsolateral, STG: superior temporal gyrus, ITG: inferior temporal gyrus, Caud/Acc: caudate-accumbens, SM: sensorimotor, V1: primary visual, A1: primary auditory, SMA: supplementary motor cortex, ASD: autism spectrum disorder, TDC: typically developing control.

Table 3: The top ten largest changes in average attractor-state activity between autistic and control individuals. Mean attractor-state activity changes are presented in the order of their absolute effect size. Reported effect sizes are mean attractor activation differences. Note that activation time series were standard scaled independently for each region, so effect size can be interpreted as showing the differences as a proportion of regional variability. All p-values are based on permutation tests (shuffling the group assignment) and corrected for multiple comparisons (via Bonferroni's correction). For a comprehensive list of significant findings, see Supplementary Figure 11.

region	*attractor*	*effect size*	*p-value*
primary auditory cortex	perception	-0.126	<0.0001
middle cingulate cortex	action	0.109	<0.0001
cerebellum lobule VIIb (medial part)	internal context	0.104	<0.0001
mediolateral sensorimotor cortex	perception	-0.099	0.00976

precuneus	action	0.098	<0.0001
middle superior temporal gyrus	perception	-0.098	<0.0001
frontal eye field	perception	-0.095	<0.0001
dorsolateral sensorimotor cortex	perception	-0.094	0.00976
posterior cingulate cortex	action	0.092	<0.0001
dorsolateral prefrontal cortex	external context	-0.092	<0.0001

Thus, we contrasted the characteristic trajectories derived from the fcANN models of the two groups (initialized with the group-level functional connectomes). Our fcANN-based flow analysis predicted that in ASD, there is an increased likelihood of states returning towards the middle (more noisy states) from the internal-external axis and an increased likelihood of states transitioning towards the extremes of the action-perception axis (Figure 5C). We observed a highly similar pattern in the real data (Pearson's r = 0.66), statistically significant after two-sided permutation testing (shuffling the group assignment; 1,000 permutations; p=0.009).

Discussion

The notion that the brain functions as an attractor network has long been proposed (Freeman, 1987; Amit, 1989; Deco & Jirsa, 2012; Deco *et al.*, 2012; Golos *et al.*, 2015; Hansen *et al.*, 2015; Vohryzek *et al.*, 2020), although the exact functional form of the network underlying large-scale brain dynamics - or at least useful approximation thereof - remained elusive. The theoretical framework of free energy minimizing attractor neural networks (FEP-ANN) (Spisak & Friston, 2025) identifies a specific class of attractor networks that emerge from first principles of self-organization - as articulated by the Free Energy Principle (FEP) (Friston, 2010; Friston *et al.*, 2023). Therefore it provides a plausible candidate model for large-scale brain attractor dynamics and yields testable predictions - measurable signatures of these special attractor networks that can be validated empirically. In this study, we have introduced, and performed initial validation of, a simple and robust network-level generative computational model, rooted in the FEP-ANN framework and providing the opportunity to test these predictions empirically. Our model, termed a functional connectivity-based attractor network (fcANN), exploits special characteristics of the emergent inference rule of FEP-ANNs. This is a stochastic rule that governs how activity evolves in time with a given set of fixed coupling weights and leads to a Markov chain Monte Carlo (MCMC) sampling process. As a consequence, the activation time series data measured in each network node can be used to reconstruct the attractor network's internal structure. Specifically, the coupling weights can be estimated as the negative inverse covariance matrix of the activation time series data. fcANN modeling applies this concept to large-scale brain dynamics as measured by resting-state fMRI data (as an estimate of weights corresponding to the steady-state distribution of the system).

The core idea underlying the fcANN reconstruction approach - the use of functional connectivity as a proxy for weighted information flow in the brain - is in line with previous empirical observations about the relationship between functional connectivity and brain activity, as articulated by the activity flow principle, first introduced by Cole *et al.* (2016). The activity flow principle states that activity in a brain region can be predicted by a weighted combination of the activity of all other regions, where the weights are set to the functional connectivity of those regions to the held-out region. This principle has been shown to hold across a wide range of experimental and clinical conditions (Cole *et al.*, 2016; Ito *et al.*, 2017; Mill *et al.*, 2022; Hearne *et al.*, 2021; Chen *et al.*, 2018). Considering that the repeated, iterative application of the activity flow equation (extended with an arbitrary sigmoidal activation function) naturally reproduces certain types of recurrent artificial neural networks, e.g. Hopfield networks (Hopfield, 1982), yields an intuitive understanding of how the fcANN model works.

However, beyond this analogy, we need concrete evidence that the fcANN model and the underlying FEP-ANN framework is a valid model of large-scale brain dynamics.

Here we have tested multiple predictions of the FEP-ANN framework. Most importantly FEP-ANNs - through their emergent predictive coding-based learning rule - have been shown to develop approximately orthogonal attractor representations, and thereby approximate the so-called Kanter-Sompolinsky projector neural networks (K-S network, for short) (Kanter & Sompolinsky, 1987). K-S networks are a special class of attractor networks that have been shown to be highly effective for pattern recognition and learning. In these networks, the attractor states are orthogonal to each other, and become equivalent with those eigenvectors of the coupling matrix, that have positive eigenvalues. This is a very strong prediction: K-S networks are a very special class of attractor networks, which do not arise from conventional learning rules, like Hebbian learning. To date, the predictive coding-based learning rule of FEP-ANNs is the only known local, incremental learning rule that can effectively approximate K-S networks in a single phase (but see "dreaming neural networks" (Fachechi et al., 2019) for a similar, two-phase approach). Thus, showing that fcANN models approximate K-S networks can be interpreted as evidence for the brain functioning akin to a FEP-ANN. Our results show that this is indeed the case: the fcANN

models reconstructed from resting-state fMRI data approximate K-S networks, and thereby exhibit approximately orthogonal attractor states.

In the FEP-ANN framework, approximate attractor orthogonality has important computational implications: it allows the system achieve maximal "storage capacity" (the number of attractors that can be stored and retrieved without interference). Furthermore, in the FEP-ANN framework, attractor states can be interpreted as learned priors that capture the statistical structure of the environment, while the stochastic dynamics implement posterior sampling. Orthogonal attractors emerge as an efficient way to span the subspace of interest, fostering generalization to unseen data (as long as it is from the subspace spanned by the existing attractors).

Next, we have demonstrated that fcANN models exhibit multiple biologically plausible attractor states, with large basins and showed that the relaxation dynamics of the reconstructed model display fast convergence to attractor states - another signature that the functional connectome being a valid proxy for the coupling weights of an underlying attractor network. Relying on previous work, we can establish a relatively straightforward (although somewhat speculative) correspondence between attractor states and brain function, mapping brain activation on the axes of internal vs. external context (Golland *et al.*, 2008; Cioli *et al.*, 2014), as well as perception vs. action (Fuster, 2004). In our framework, the attractor states can be interpreted as learned priors that capture the statistical structure of the environment, while the stochastic dynamics implement posterior sampling. This connection suggests that canonical resting-state networks may represent the brain's internal generative model of the world, continuously updated through the emergent learning dynamics we described theoretically. Furthermore, the relation between fcANN models and the FEP-ANN framework substantiates that the reconstructed attractor states are not solely local minima in the state space but act as a driving force for the dynamic trajectories of brain activity. We argue that attractor dynamics may be the main driving factor for the spatial and temporal autocorrelation structure of the brain, recently described to be predictive of network topology in relation to age, subclinical symptoms of dementia, and pharmacological manipulations with serotonergic drugs (Shinn *et al.*, 2023). Nevertheless, attractor states should not be confused with the conventional notion of brain substates (Chen *et al.*, 2015; Kringelbach & Deco, 2020) and large-scale functional gradients (Margulies *et al.*, 2016). In the fcANN framework, attractor states can rather be conceptualized as "Platonic idealizations" of brain activity, that are continuously approximated - but never reached - by the brain, resulting in re-occurring patterns (brain substates) and smooth gradual transitions (large-scale gradients).

Considering the functional connectome as weights of a neural network distinguishes our methodology from conventional biophysical and phenomenological computational modeling strategies, which usually rely on the structural connectome to model polysynaptic connectivity (Cabral *et al.*, 2017; Deco *et al.*, 2012; Golos *et al.*, 2015; Hansen *et al.*, 2015). Given the challenges of accurately modeling the structure-function coupling in the brain (Seguin *et al.*, 2023), such models are currently limited in terms of reconstruction accuracy, hindering translational applications. By working with direct, functional MRI-based activity flow estimates, fcANNs bypass the challenge of modeling the structural-functional coupling and are able to provide a more accurate representation of the brain's dynamic activity propagation (although at the cost of losing the ability to provide biophysical details on the underlying mechanisms). Another advantage of the proposed model is its simplicity. While many conventional computational models rely on the optimization of a high number of free parameters, the basic form of the fcANN approach comprises solely two, easily interpretable "hyperparameters" (temperature and noise) and yields notably consistent outcomes across an extensive range of these parameters (Supplementary Figure 3, 5, 7, 8, 9). To underscore the potency of this simplicity and stability, in the present work, we avoided any unnecessary parameter optimization, leaving a negligible chance of overfitting. It is likely, however, that extensive parameter optimization could further improve the accuracy of the model.

Further, the fcANN approach links brain dynamics directly to dynamical systems theory and the free energy principle, conceptualizes the emergence of large-scale canonical brain networks (Zalesky et al., 2014) in terms of multistability, and sheds light on the origin of characteristic task-responses that are accounted for by "ghost attractors" in the system (Deco & Jirsa, 2012; Vohryzek *et al.*, 2020). As fcANNs do not need to be trained to solve any explicit tasks, they are well suited to examine spontaneous brain dynamics. However, it is worth mentioning that fcANNs can also be further trained via the predictive coding-based learning rule of FEP-ANNs, to solve various tasks or to match developmental dynamics or pathological alterations. In this promising future direction, the training procedure itself becomes part of the model, providing testable hypotheses about the formation, and various malformations, of brain dynamics. A promising application of this is to consider structural brain connectivity (as measured by diffusion MRI) as a sparsity constraint for the coupling weights and then train the fcANN model to match the observed resting-state brain dynamics. If the resulting structural-functional ANN model is able to closely match the observed functional brain substate dynamics, it can be used as a novel approach to quantify and understand the structural functional coupling in the brain.

Given its simplicity, it is noteworthy how well the fcANN model is able to reconstruct and predict brain dynamics under a wide range of conditions. First and foremost, we have found that the topology of the functional

connectome seems to be well suited to function as an attractor network, as it converges much faster than the respective null models. Second, we found that the two-dimensional fcANN projection can explain more variance in real (unseen) resting-state fMRI data than the first two principal components derived from the data itself. This may indicate that through the known noise tolerance of attractor neural networks, fcANNs are able to capture essential characteristics of the underlying dynamic processes even if our empirical measurements are corrupted by noise and low sampling rate. Indeed, fcANN attractor states were found to be robust to noisy weights (Supplementary Figure 10) and highly replicable across datasets acquired at different sites, with different scanners and imaging sequences (study 2 and 3). The observed high level of replicability allowed us to re-use the fcANN model constructed with the functional connectome of study 1 for all subsequent analyses, without any further fine-tuning or study-specific parameter optimization.

Both conceptually and in terms of analysis practices, resting and task states are often treated as separate phenomena. However, in the fcANN framework, the differentiation between task and resting states is considered an artificial dichotomy. Task-based brain activity in the fcANN framework is not a mere response to external stimuli in certain brain locations but a perturbation of the brain's characteristic dynamic trajectories, with increased preference for certain locations on the energy landscape ("ghost attractors"). In our analyses, the fcANN approach captured and predicted participant-level activity changes induced by pain and its self-regulation and gave a mechanistic account for how relatively small activity changes in a single region (NAcc) may result in a significantly altered pain experience. Our control-signal analysis is different from, but compatible with, linear network control theory-based approaches (Liu *et al.*, 2011; Gu *et al.*, 2015). Combining network control theory with the fcANN approach could provide a powerful framework for understanding the effects of various tasks, conditions, and interventions (e.g., brain stimulation) on brain dynamics.

Brain dynamics can not only be perturbed by task or other types of experimental or naturalistic interventions, but also by pathological alterations. Here we provide an initial demonstration (study 7) of how fcANN-based analyses can characterize and predict altered brain dynamics in autism spectrum disorder (ASD). The observed ASD-associated changes in brain dynamics are indicative of a reduced ability to flexibly switch between perception and internal representations, corroborating previous findings that in ASD, sensory-driven connectivity transitions do not converge to transmodal areas (Hong *et al.*, 2019). Such findings are in line with previous reports of a reduced influence of context on the interpretation of incoming sensory information in ASD (e.g., the violation of Weber's law) (Hadad & Schwartz, 2019).

Our findings open up a series of exciting opportunities for the better understanding of brain function in health and disease. First, fcANN analyses may provide insights into the causes of changes in brain dynamics, by, for instance, identifying the regions or connections that act as an "Achilles' heel" in generating such changes. Such control analyses could, for instance, aid the differentiation of primary causes and secondary effects of activity or connectivity changes in various clinical conditions. Rather than viewing pathology as static connectivity differences, our approach suggests that disorders may reflect altered attractor landscapes that bias brain dynamics toward maladaptive states. This perspective could inform the development of targeted interventions that aim to reshape these landscapes through neurofeedback, brain stimulation, or pharmacological approaches.

Second, as a generative model, fcANNs provide testable predictions about the effects of various interventions on brain dynamics, including pharmacological modulations as well as non-invasive brain stimulation (e.g., transcranial magnetic or direct current stimulation, focused ultrasound, etc.) and neurofeedback. Obtaining the optimal stimulation or treatment target within the fcANN framework (e.g., by means of network control theory (Liu *et al.*, 2011)) is one of the most promising future directions with the potential to significantly advance the development of novel, personalized treatment approaches.

Third, the theoretical integration of the fcANN model with the FEP-ANN framework positions our work within a broader scientific program that seeks to understand the brain as a self-organizing, information-processing system governed by fundamental physical and computational principles. The empirical validation of attractor orthogonality represents a crucial step toward establishing this unified framework for understanding brain function across scales and contexts.

The proposed approach is not without limitations. First, fcANNs do not incorporate information about anatomical connectivity and do not explicitly model biophysical details. Thus, in its present form, the model is not suitable to study the structure-function coupling and cannot yield mechanistic explanations underlying (altered) polysynaptic connections, at the level of biophysical details. Nevertheless, our approach showcases that many characteristics of brain dynamics, like multistability, temporal autocorrelations, states and gradients, can be explained, and predicted, by a very simple nonlinear phenomenological model. Second, our model assumes a stationary functional connectome, which seems to contradict notions of dynamic connectivity. However, while the underlying FEP-ANN framework focuses on the long-term steady-state distribution of the system, it also naturally incorporates multistable fluctuations and the related dynamic connectivity changes through the

stochastic relaxation dynamics. This is in line with the notion of "latent functional connectivity", an intrinsic brain network architecture built up from connectivity properties that are persistent across brain substates (McCormick *et al.*, 2022).

In this initial work, we presented the simplest possible implementation of the fcANN concept. It is clear that the presented analyses exploit only a small proportion of the richness of the full state-space dynamics reconstructed by the fcANN model. There are many potential ways to further improve the utility of the fcANN approach. Increasing the number of reconstructed attractor states (by increasing the temperature parameter), investigating higher-dimensional dynamics, fine-tuning the hyperparameters, and testing the effect of different initializations and perturbations are all important directions for future work, with the potential to further improve the model's accuracy and usefulness.

Conclusion

Here we have proposed a principled, lightweight, theory-driven framework that instantiates the inference dynamics of free-energy-minimizing attractor networks (FEP-ANNs) to model large-scale brain activity. Initialized with empirical functional connectivity, the fcANN links brain connectivity to activity and identifies neurobiologically meaningful attractor states underlying large-scale brain dynamics. We demonstrated that the fcANN models display signs of attractor self-orthogonalization - a hallmark of FEP-ANN systems. The proposed framework provides a simple, interpretable, and predictive basis for studying rest, task perturbations, and disease, and for model-guided interventions.

Acknowledgements

The work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation; projects 'TRR289 - Treatment Expectation', ID 422744262 and 'SFB1280 - Extinction Learning', ID 316803389) and by IBS-R015-D1 (Institute for Basic Science; C.W.-W.).

Analysis source code

https://github.com/pni-lab/connattractor

Project website

https://pni-lab.github.io/connattractor/

Data availability

Studies 1, 2 and 4 are available at openneuro.org (ds002608, ds002608, ds000140). Data for study 3 are available upon request. Data for studies 5-6 are available at the GitHub page of the project: https://github.com/pni-lab/connattractor. Study 7 is available at https://fcon_1000.projects.nitrc.org/indi/abide/, preprocessed data are available at http://preprocessed-connectomes-project.org/.

Declaration of interests

The authors declare no competing interests.

References

Amit, D. J. (1989). *Modeling Brain Function: The World of Attractor Neural Networks*. Cambridge University Press. 10.1017/cbo9780511623257

Barttfeld, P., Uhrig, L., Sitt, J. D., Sigman, M., Jarraya, B., & Dehaene, S. (2015). Signature of consciousness in the dynamics of resting-state brain activity. *Proceedings of the National Academy of Sciences*, *112*(3), 887–892.

Bassett, D. S., & Sporns, O. (2017). Network neuroscience. *Nature Neuroscience*, *20*(3), 353–364.

Breakspear, M. (2017). Dynamic models of large-scale brain activity. *Nature Neuroscience*, *20*(3), 340–352.

Buzsaki, G. (2006). *Rhythms of the Brain*. Oxford university press.

Cabral, J., Kringelbach, M. L., & Deco, G. (2017). Functional connectivity dynamically evolves on multiple time-scales over a static structural connectome: Models and mechanisms. *NeuroImage*, *160*, 84–96.

Chen, J. E., Chang, C., Greicius, M. D., & Glover, G. H. (2015). Introducing co-activation pattern metrics to quantify spontaneous brain network dynamics. *Neuroimage*, *111*, 476–488.

Chen, R. H., Ito, T., Kulkarni, K. R., & Cole, M. W. (2018). The human brain traverses a common activation-pattern state space across task and rest. *Brain Connectivity*, *8*(7), 429–443.

Chiêm, B., Crevecoeur, F., & Delvenne, J.-C. (2021). Structure-informed functional connectivity driven by identifiable and state-specific control regions. *Network Neuroscience*, *5*(2), 591–613.

Cioli, C., Abdi, H., Beaton, D., Burnod, Y., & Mesmoudi, S. (2014). Differences in human cortical gene expression match the temporal properties of large-scale functional networks. *PloS One*, *9*(12), e115913.

Cole, M. W. (2024). *The explanatory power of activity flow models of brain function*. arXiv. 10.48550/ARXIV.2402.02191

Cole, M. W., Ito, T., Bassett, D. S., & Schultz, D. H. (2016). Activity flow over resting-state networks shapes cognitive task activations. *Nature Neuroscience*, *19*(12), 1718–1726.

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

Deco, G., & Jirsa, V. K. (2012). Ongoing cortical activity at rest: criticality, multistability, and ghost attractors. *Journal of Neuroscience*, *32*(10), 3366–3375.

Deco, G., Senden, M., & Jirsa, V. (2012). How anatomy shapes dynamics: a semi-analytical study of the brain at rest by a simple spin model. *Frontiers in Computational Neuroscience*, *6*, 68.

Di Martino, A., Yan, C.-G., Li, Q., Denio, E., Castellanos, F. X., Alaerts, K., Anderson, J. S., Assaf, M., Bookheimer, S. Y., Dapretto, M., & others. (2014). The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular Psychiatry*, *19*(6), 659–667.

Dosenbach, N. U. F., Raichle, M. E., & Gordon, E. M. (2025). The brain's action-mode network. *Nature Reviews Neuroscience*, *26*(3), 158–168. 10.1038/s41583-024-00895-x

Dotsenko, V. S., & Tirozzi, B. (1991). Replica symmetry breaking in neural networks with modified pseudo-inverse interactions. *Journal of Physics A: Mathematical and General*, *24*(21), 5163–5180. 10.1088/0305-4470/24/21/026

Fachechi, A., Agliari, E., & Barra, A. (2019). Dreaming neural networks: Forgetting spurious memories and reinforcing pure ones. *Neural Networks*, *112*, 24–40. 10.1016/j.neunet.2019.01.006

Freeman, W. J. (1987). Simulation of chaotic EEG patterns with a dynamic model of the olfactory system. *Biological Cybernetics*, *56*(2–3), 139–150.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138. 10.1038/nrn2787

Friston, K., Da Costa, L., Sajid, N., Heins, C., Ueltzhöffer, K., Pavliotis, G. A., & Parr, T. (2023). The free energy principle made simpler but not too simple. *Physics Reports*, *1024*, 1–29. 10.1016/j.physrep.2023.07.001

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O'Doherty, J., & Pezzulo, G. (2016). Active inference and learning. *Neuroscience & Biobehavioral Reviews*, *68*, 862–879. 10.1016/j.neubiorev.2016.06.022

Fuster, J. M. (2004). Upper processing stages of the perception–action cycle. *Trends in Cognitive Sciences*, *8*(4), 143–145.

Golland, Y., Golland, P., Bentin, S., & Malach, R. (2008). Data-driven clustering reveals a fundamental subdivision of the human cortex into two global systems. *Neuropsychologia*, *46*(2), 540–553.

Golos, M., Jirsa, V., & Daucé, E. (2015). Multistability in large scale models of brain activity. *PLoS Computational Biology*, *11*(12), e1004644.

Greene, A. S., Horien, C., Barson, D., Scheinost, D., & Constable, R. T. (2023). Why is everyone talking about brain state? *Trends in Neurosciences*.

Gu, S., Pasqualetti, F., Cieslak, M., Telesford, Q. K., Yu, A. B., Kahn, A. E., Medaglia, J. D., Vettel, J. M., Miller, M. B., Grafton, S. T., & others. (2015). Controllability of structural brain networks. *Nature Communications*, *6*(1), 8414.

Hadad, B.-S., & Schwartz, S. (2019). Perception in autism does not adhere to Weber's law. *Elife*, *8*, e42223.

Hansen, E. C., Battaglia, D., Spiegler, A., Deco, G., & Jirsa, V. K. (2015). Functional connectivity dynamics: modeling the switching behavior of the resting state. *Neuroimage*, *105*, 525–535.

Hearne, L. J., Mill, R. D., Keane, B. P., Repovš, G., Anticevic, A., & Cole, M. W. (2021). Activity flow underlying abnormalities in brain activations and cognition in schizophrenia. *Science Advances*, *7*(29), eabf2513.

Heinz, A., Murray, G. K., Schlagenhauf, F., Sterzer, P., Grace, A. A., & Waltz, J. A. (2019). Towards a unifying cognitive, neurophysiological, and computational neuroscience account of schizophrenia. *Schizophrenia Bulletin*, *45*(5), 1092–1100.

Hong, S.-J., Vos de Wael, R., Bethlehem, R. A., Lariviere, S., Paquola, C., Valk, S. L., Milham, M. P., Di Martino, A., Margulies, D. S., Smallwood, J., & others. (2019). Atypical functional connectome hierarchy in autism. *Nature Communications*, *10*(1), 1022.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, *79*(8), 2554–2558.

Hopfield, J. J., Feinstein, D. I., & Palmer, R. G. (1983). 'Unlearning' has a stabilizing effect in collective memories. *Nature*, *304*(5922), 158–159. 10.1038/304158a0

Huntenburg, J. M., Bazin, P.-L., & Margulies, D. S. (2018). Large-scale gradients in human cortical organization. *Trends in Cognitive Sciences*, *22*(1), 21–31.

Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., Della Penna, S., Duyn, J. H., Glover, G. H., Gonzalez-Castillo, J., & others. (2013). Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage*, *80*, 360–378.

Ito, T., Kulkarni, K. R., Schultz, D. H., Mill, R. D., Chen, R. H., Solomyak, L. I., & Cole, M. W. (2017). Cognitive task information is transferred between brain regions via resting-state network topology. *Nature Communications*, *8*(1), 1027.

Kanter, I., & Sompolinsky, H. (1987). Associative recall of memory without errors. *Physical Review A*, *35*(1), 380–392. 10.1103/physreva.35.380

Khona, M., & Fiete, I. R. (2022). Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, *23*(12), 744–766.

Koiran, P. (1994). Dynamics of discrete time, continuous state Hopfield networks. *Neural Computation*, *6*(3), 459–468.

Kriegeskorte, N., & Douglas, P. K. (2018). Cognitive computational neuroscience. *Nature Neuroscience*, *21*(9), 1148–1160.

Kringelbach, M. L., & Deco, G. (2020). Brain States and Transitions: Insights from Computational Neuroscience. *Cell Reports*, *32*(10), 108128. 10.1016/j.celrep.2020.108128

Liu, X., & Duyn, J. H. (2013). Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proceedings of the National Academy of Sciences*, *110*(11), 4392–4397.

Liu, Y.-Y., Slotine, J.-J., & Barabási, A.-L. (2011). Controllability of complex networks. *Nature*, *473*(7346), 167–173.

Margulies, D. S., Ghosh, S. S., Goulas, A., Falkiewicz, M., Huntenburg, J. M., Langs, G., Bezgin, G., Eickhoff, S. B., Castellanos, F. X., Petrides, M., & others. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences*, *113*(44), 12574–12579.

McCormick, E. M., Arnemann, K. L., Ito, T., Hanson, S. J., & Cole, M. W. (2022). Latent functional connectivity underlying multiple brain states. *Network Neuroscience*, *6*(2), 570–590. 10.1162/netn_a_00234

Medrano, J., Friston, K., & Zeidman, P. (2024). Linking fast and slow: The case for generative models. *Network Neuroscience*, *8*(1), 24–43. 10.1162/netn_a_00343

Meer, J. N. van der, Breakspear, M., Chang, L. J., Sonkusare, S., & Cocchi, L. (2020). Movie viewing elicits rich and reliable brain state dynamics. *Nature Communications*, *11*(1), 5004.

Mill, R. D., Hamilton, J. L., Winfield, E. C., Lalta, N., Chen, R. H., & Cole, M. W. (2022). Network modeling of dynamic brain interactions predicts emergence of neural information that supports human cognitive behavior. *PLoS Biology*, *20*(8), e3001686.

Millidge, B., Song, Y., Salvatori, T., Lukasiewicz, T., & Bogacz, R. (2022). *A Theoretical Framework for Inference and Learning in Predictive Coding Networks*. arXiv. 10.48550/ARXIV.2207.12316

Murray, J. D., Demirtaş, M., & Anticevic, A. (2018). Biophysical modeling of large-scale brain dynamics and applications for computational psychiatry. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, *3*(9), 777–787.

Padmanabhan, A., Lynch, C. J., Schaer, M., & Menon, V. (2017). The default mode network in autism. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, *2*(6), 476–486.

Papadopoulos, L., Kim, J. Z., Kurths, J., & Bassett, D. S. (2017). Development of structural correlations and synchronization from adaptive rewiring in networks of Kuramoto oscillators. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *27*(7).

Plakhov, A. Yu. (n.d.). The converging unlearning algorithm for the Hopfield neural network: optimal strategy. *Proceedings of the 12th IAPR International Conference on Pattern Recognition (Cat. No.94CH3440-5)*, *2*, 104–106. 10.1109/icpr.1994.576884

Poerio, G. L., & Karapanagiotidis, T. (2025). The default mode network and the complex dynamics of ongoing experience: an attractor-state perspective. *Current Opinion in Behavioral Sciences*, *65*, 101546. 10.1016/j.cobeha.2025.101546

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, *98*(2), 676–682.

Sanchez-Romero, R., Ito, T., Mill, R. D., Hanson, S. J., & Cole, M. W. (2023). Causally informed activity flow models provide mechanistic insight into network-generated cognitive activations. *NeuroImage*, *278*, 120300. 10.1016/j.neuroimage.2023.120300

Scheid, B. H., Ashourvan, A., Stiso, J., Davis, K. A., Mikhail, F., Pasqualetti, F., Litt, B., & Bassett, D. S. (2021). Time-evolving controllability of effective connectivity networks during seizure progression. *Proceedings of the National Academy of Sciences*, *118*(5), e2006436118.

Schiff, S. J., Jerger, K., Duong, D. H., Chang, T., Spano, M. L., & Ditto, W. L. (1994). Controlling chaos in the brain. *Nature*, *370*(6491), 615–620.

Schirner, M., Kong, X., Yeo, B. T., Deco, G., & Ritter, P. (2022). Dynamic primitives of brain network interaction. *NeuroImage*, *250*, 118928.

Seguin, C., Sporns, O., & Zalesky, A. (2023). Brain network communication: concepts, models and applications. *Nature Reviews Neuroscience*, *24*(9), 557–574.

Shinn, M., Hu, A., Turner, L., Noble, S., Preller, K. H., Ji, J. L., Moujaes, F., Achard, S., Scheinost, D., Constable, R. T., & others. (2023). Functional brain networks reflect spatial and temporal autocorrelation. *Nature Neuroscience*, 1–12.

Smith, S. M., Miller, K. L., Moeller, S., Xu, J., Auerbach, E. J., Woolrich, M. W., Beckmann, C. F., Jenkinson, M., Andersson, J., Glasser, M. F., & others. (2012). Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences*, *109*(8), 3131–3136.

Spisak, T., & Friston, K. (2025). *Self-orthogonalizing attractor neural networks emerging from the free energy principle*. arXiv. 10.48550/ARXIV.2505.22749

Sydnor, V. J., Larsen, B., Bassett, D. S., Alexander-Bloch, A., Fair, D. A., Liston, C., Mackey, A. P., Milham, M. P., Pines, A., Roalf, D. R., & others. (2021). Neurodevelopment of the association cortices: Patterns, mechanisms, and implications for psychopathology. *Neuron*, *109*(18), 2820–2846.

Vidaurre, D., Smith, S. M., & Woolrich, M. W. (2017). Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, *114*(48), 12827–12832.

Vohryzek, J., Deco, G., Cessac, B., Kringelbach, M. L., & Cabral, J. (2020). Ghost attractors in spontaneous brain activity: Recurrent excursions into functionally-relevant BOLD phase-locking states. *Frontiers in Systems Neuroscience*, *14*, 20.

Wagner, D. H. (1989). The Existence and Behavior of Viscous Structure for Plane Detonation Waves. *SIAM Journal on Mathematical Analysis*, *20*(5), 1035–1054. 10.1137/0520069

Woo, C.-W., Roy, M., Buhle, J. T., & Wager, T. D. (2015). Distinct brain systems mediate the effects of nociceptive input and self-regulation on pain. *PLoS Biology*, *13*(1), e1002036.

Zalesky, A., Fornito, A., Cocchi, L., Gollo, L. L., & Breakspear, M. (2014). Time-resolved resting-state brain networks. *Proceedings of the National Academy of Sciences*, *111*(28), 10341–10346.

Zunhammer, M., Spisák, T., Wager, T. D., & Bingel, U. (2021). Meta-analysis of neural systems underlying placebo analgesia from individual participant fMRI data. *Nature Communications*, *12*(1), 1391.

Methods

Statistical analysis and reporting

Unless stated otherwise, all hypothesis tests are two-sided. When permutation tests are used, reported p-values are empirical permutation p-values computed as the proportion of null replicates with a test statistic at least as extreme as the observed one (with the usual +1 numerator/denominator correction). We report the statistic used in each comparison (e.g., Pearson correlation r, chi-square dissimilarity, L2 norm of position differences, median iterations to convergence). Degrees of freedom are reported where parametric tests are used; permutation-based p-values do not have degrees of freedom. By default, permutation counts were 1,000, except where explicitly noted (e.g., 50,000 label shuffles in clinical analyses). Confidence intervals based on bootstrap are percentile bootstrap intervals; where shown as 99% CIs, they reflect the percentile range from the bootstrap samples described below. For multiple comparisons in the clinical analyses, we applied Bonferroni correction across regions and attractor states as specified below.

Data

We obtained functional MRI data from 7 sources (Methods Table 1). MRI sequence parameters for studies 1–4 are summarized in Supplementary Table 1. We included three resting-state studies with healthy volunteers (study 1, study 2, study 3, $n_{total} = 118$), one task-based study (study 4, $n_{total} = 33$ participants, 9 runs each), an individual-participant meta-analytic activation map of pain (study 5, $n_{total} = 603$ from 20 studies), 8 task-based activation patterns obtained from coordinate-based meta-analyses via Neurosynth (study 6, 14,371 studies in total; see Supplementary Table 2), and a resting-state dataset focusing on autism spectrum disorder (ASD) from ABIDE (study 7, $n_{total} = 1{,}112$, Di Martino *et al.* (2014)).

Methods Table 1: Datasets and studies. The table includes details about the study modality, analysis aims, sample size used for analyses, mean age, gender ratio, and references.

study	*modality*	*analysis*	*n*	*age (mean±sd)*	*%female*	*references*
study 1	resting state	discovery	41	26.1±3.9	37%	Spisak *et al.* (2020)
study 2	resting state	replication	48	24.9±3.5	54%	Spisak *et al.* (2020)
study 3	resting state	replication	29	24.8±3.1	53%	Spisak *et al.* (2020)
study 4	task-based	pain self-regulation	33	27.9 ± 9.0	66%	Woo *et al.* (2015)
study 5 (Metaanalysis)	task-based	IPD meta-analysis pain map	n=603 (20 studies)	26.3 ± 5.9	39%	Zunhammer *et al.* (2021)
study 6 (Neurosynth)	task-based	coordinate-based meta-analyses	14371 studies in total	N/A	N/A	Yarkoni *et al.* (2011)
study 7 (ABIDE, NYU sample)	resting state	Autism Spectrum Disorder	ASD: 98; NC: 74	15.3±6.6	20.9%	Di Martino *et al.* (2014)

Study 1 was used to evaluate whether the resting-state functional connectome can be treated as an attractor network, to optimize the temperature (β) and noise (ϵ) parameters of the fcANN model, and to evaluate the proposed approach for reconstructing resting-state brain dynamics. Studies 2 and 3 served as replication datasets. Studies 1–3 are well suited to examine replicability and generalizability; data were acquired in three centers across two countries, by different research staff, with different scanners (Philips, Siemens, GE) and imaging sequences. Further details on studies 1–3 are described in Spisak *et al.* (2020). The ability of the proposed approach to model task-based perturbations of brain dynamics was evaluated in study 4, which consisted of nine task-based fMRI runs for each of the 33 healthy volunteers. In all runs, participants received heat pain stimulation. Each stimulus lasted 12.5 seconds, with a 3-second ramp-up and 2-second ramp-down and 7.5 seconds at target temperature. Six temperature levels were administered (44.3°C, 45.3°C, 46.3°C, 47.3°C, 48.3°C, 49.3°C). We used run 1 (passive experience), run 3 (down-regulation), and run 7 (up-regulation). In runs 3 and 7, participants were asked to cognitively increase (regulate up) or decrease (regulate down) pain intensity. No self-regulation instructions were

provided in run 1. See Woo *et al.* (2015) for details. Pain control signal for our task-based trajectory analyses on data from study 4 was derived from our individual participant meta-analysis of 20 pain fMRI studies (study 5, n=603). For details, see Zunhammer *et al.* (2021). To obtain fMRI activation maps for other tasks, we used Neurosynth (Tor D. (2011)), a web-based platform for large-scale, automated synthesis of fMRI data. We performed eight coordinate-based meta-analyses with the terms "motor", "auditory", "visual", "face", "autobiographical", "theory mind", "language", and "pain" (Supplementary Table 2) and obtained Z-score maps from a two-way ANOVA, comparing coordinates reported for studies with and without the term of interest and testing for a nonzero association between term use and voxel activation. In study 7 (ABIDE), we obtained preprocessed regional time-series data from the Preprocessed Connectome Project (Craddock *et al.*, 2013), as shared at https://osf.io/hc4md by Dadi *et al.* (2019). Preprocessed time-series data were obtained with the 122-region version of the BASC (Bootstrap Analysis of Stable Clusters) brain atlas (Bellec *et al.*, 2010).

Preprocessing and time-series extraction

Functional MRI data from studies 1–4 were preprocessed with our in-house analysis pipeline, the RPN-pipeline (https://github.com/spisakt/RPN-signature). The RPN-pipeline is based on PUMI (Neuroimaging Pipelines Using Modular workflow Integration, https://github.com/pni-lab/PUMI), a nipype-based (Gorgolewski *et al.*, 2011) workflow management system. It capitalizes on tools from FSL (Jenkinson *et al.*, 2012), ANTs (Avants *et al.*, 2011), and AFNI (Cox, 1996), with code partially adapted from C-PAC (Craddock *et al.*, 2013) and niworkflows (Esteban *et al.*, 2019), as well as in-house Python routines.

Brain extraction from both anatomical and structural images, as well as tissue segmentation from the anatomical images, was performed with FSL BET and FAST. Anatomical images were linearly and nonlinearly co-registered to the 1 mm MNI152 standard brain template with ANTs (see https://gist.github.com/spisakt/0caa7ec4bc18d3ed736d3a4e49da7415 for parameters). Functional images were co-registered to the anatomical images with FSL FLIRT's boundary-based registration. All resulting transformations were saved for further use. Preprocessing of functional images was performed in native space, without resampling. Realignment-based motion correction was performed with FSL MCFLIRT. The resulting six head motion estimates (three rotations, three translations), their squared terms, their derivatives, and the squared derivatives (Friston-24 expansion, Friston *et al.* (1996)) were calculated as nuisance signals. Additionally, head motion was summarized as framewise displacement (FD) time series, according to Power's method (Power *et al.*, 2012), for use in censoring and exclusion. After motion correction, outliers (e.g., motion spikes) in time-series data were attenuated using AFNI despike. The union of eroded white-matter maps and ventricle masks was transformed to native functional space and used to extract noise signals for anatomical CompCor correction (Behzadi *et al.*, 2007).

In a nuisance regression step, six CompCor parameters (the first six principal components of noise-region time series), the Friston-24 motion parameters, and the linear trend were removed with a general linear model. On the residual data, temporal band-pass filtering was performed with AFNI's 3dBandpass to retain the 0.008–0.08 Hz band. To further attenuate motion artifacts, potentially motion-contaminated time frames, defined by a conservative FD > 0.15 mm threshold, were dropped (scrubbing, Satterthwaite *et al.* (2013)). Participants were excluded if more than 50% of frames were scrubbed.

The 122-parcel BASC (Bootstrap Analysis of Stable Clusters) atlas (Bellec *et al.*, 2010) was individualized by transforming it to each participant's native functional space (nearest-neighbour interpolation) and masking by that participant's grey-matter mask to retain only grey-matter voxels. Voxel time series were then averaged within BASC regions.

All these preprocessing steps are part of the containerized version of the RPN-pipeline (https://spisakt.github.io/RPN-signature), which we run with default parameters for all studies, as in Spisak *et al.* (2020).

Functional connectome

Regional time series were ordered into large-scale functional modules (defined by the 7-parcel level of the BASC atlas) for visualization. Next, in all datasets, we estimated study-level mean connectivity matrices by regularized partial correlation via the Ledoit–Wolf shrinkage estimator (Ledoit & Wolf, 2004), as implemented in Nilearn (Abraham *et al.*, 2014). The precision matrix was obtained by inverting the shrinkage covariance, and partial correlations were derived from the precision; diagonal elements were set to zero.

fcANN inference (FEP-ANN) and update rules

Our fcANN instantiates the inference dynamics of free-energy-minimizing attractor neural networks (FEP-ANNs) at the macro-scale. Each node represents a brain region with continuous activity $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_m)$, and couplings

are given by the symmetrized matrix $\boldsymbol{J}$ (see Functional connectome). Unless noted otherwise, biases are zero ($\boldsymbol{b} = \boldsymbol{0}$).

Deterministic inference. In the noise-free symmetric case, activities are updated by repeatedly applying a sigmoidal nonlinearity to the weighted input

$$\boldsymbol{\sigma}^{(t+1)} = S\left(\beta\, \boldsymbol{J}\, \boldsymbol{\sigma}^{(t)}\right), \tag{1}$$

where S is a smooth odd sigmoid (we used $tanh$ as a practical, fast surrogate for the Langevin function) and β is the inverse temperature (precision) scaling the couplings. As the inference rule was derived as a gradient descent on free energy, iterations monotonically decrease the free energy function and therefore converge to a local free-energy minimum without any external optimizer. Thus, convergence does not require any optimization procedure with an external optimizer. Instead, it arises as the fixed point of repeated local inference updates, which implement gradient descent on free energy in the deterministic symmetric case (see main text).

Stochastic (Langevin-style) inference. For generative modeling of dynamics, we adopt a slight variation of the FEP-ANN inference rule: starting from the deterministic update above, we add zero-mean Gaussian noise directly to the post-activation state (Langevin-style)

$$\boldsymbol{\sigma}^{(t+1)} = S\left(\beta\, \boldsymbol{J}\, \boldsymbol{\sigma}^{(t)}\right) + \boldsymbol{\omega}^{(t)}, \quad \boldsymbol{\omega}^{(t)} \sim \mathcal{N}(\boldsymbol{0}, \epsilon^2 \boldsymbol{I}). \tag{2}$$

This explicit additive Gaussian noise differs from the continuous-Bernoulli noise implied by the theoretical derivation but aligns with common Langevin formulations and was empirically robust. We use deterministic updates to identify attractors and study convergence; we use stochastic updates as a generative model of multistable dynamics.

fcANN convergence and attractors

We investigated convergence under the deterministic update (Eq. (1)) by contrasting iterations-to-convergence of the empirical fcANN against a permutation-based null. The null was constructed by randomly permuting the upper triangle of $\boldsymbol{J}$ and reflecting it to preserve symmetry (destroying topology while preserving weight distribution). For each of 1,000 permutations, we initialized both models with the same random state and counted iterations to convergence. Statistical significance of faster convergence in the empirical connectome was assessed via a one-sided Wilcoxon signed-rank test on paired iteration counts (1,000 pairs), testing whether the empirical connectome converges in fewer iterations than its permuted counterpart. We repeated this procedure across inverse-temperature values $\beta \in \{0.035, 0.040, 0.045, 0.050, 0.055, 0.060\}$ (yielding 2–8 attractor states). See Supplementary Figure 6 for detailed results.

fcANN projection

We mapped out the fcANN state space by initializing the model with a random input and applying the stochastic update (Eq. (2)) for iterations, storing visited activity configurations. We performed principal component analysis (PCA) on the samples and used the first two PCs as the coordinate system for the fcANN projection. Using multinomial logistic regression, we predicted attractor identity from the first two PCs. Performance was evaluated with 10-fold cross-validation and a two-sided permutation test (1,000 label permutations within folds). We visualized attractor positions and decision boundaries based on this classifier. The fcANN projection is used for visualization; unless otherwise noted, inferential comparisons between empirical and simulated data are performed in the full 122-dimensional space.

Parameter optimization

Based on the convergence analysis, we selected the inverse temperature providing the fastest median convergence ($\beta = 0.04$) for subsequent analyses, to minimize computational costs. We then optimized the stochastic noise level ϵ for the Langevin-style update by comparing the full 122-dimensional distributions of empirical and fcANN-generated data.

Specifically, for eight ϵ values spaced logarithmically between 0.1 and 1, we generated samples with Eq. (2) and computed the 2-Wasserstein (Earth Mover's) distance between the 122-dimensional empirical and simulated distributions (after per-region z-scoring). As a null, we drew surrogate samples from a covariance-matched multivariate normal distribution $\mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma}$ is the empirical covariance of the regional time series. For each ϵ, we generated 1,000 null surrogates and contrasted the empirical Wasserstein distance to the null distribution, summarizing the separation with Glass's Δ and permutation p-values. We selected $\epsilon = 0.37$ for all subsequent analyses.

Replicability

We obtained the four attractor states in study 1, as described above. We then constructed two other fcANNs, based on the study-mean functional connectome obtained in studies 2 and 3 and obtained all attractor states of these models, with the same parameter settings ($\beta = 0.04$ and $\epsilon = 0.37$) as in study 1. In both replication studies we found four attractor states. The spatial similarity of attractor states across studies was evaluated by Pearson's correlation coefficient.

Evaluation: resting state dynamics

To evaluate the explanatory power of the fcANN projection, we performed PCA on the preprocessed fMRI time frames from study 1 (analogous to the methodology of the fcANN projection, but on the empirical time-series data). Next, we fitted linear regression models using the first two fcANN-based or data-based PCs as regressors to reconstruct the real fMRI time frames. In-sample explained variances and confidence intervals were calculated for both models with bootstrapping (100 samples; percentile 99% CIs). Differences in explained variance between fcANN- and data-based PCs were assessed with a two-sided percentile bootstrap on the difference in R^2. To evaluate out-of-sample generalization of the PCs (fcANN- and data-based) from study 1, we calculated how much variance they explain in study 2.

Similarity between state occupancy and distribution was calculated during parameter optimization. More detail on the associated null models can be found in Supplementary Figure 7.

To confirm that the real and fcANN temporal sequences (from the stochastic relaxation) display similar temporal autocorrelation properties, we compared both to their randomly shuffled variant with a "flow analysis". First, we calculated the direction on the fcANN projection plane between each successive TR (a vector on the fcANN projection plane for each TR transition), both for the empirical and the shuffled data. Next, we obtained two-dimensional binned means for both the x and y coordinates of these transition vectors (pooled across all participants), calculated over a 100×100 grid of uniformly distributed bins in the [−6, 6] range (arbitrary units) and applied Gaussian smoothing with $\sigma = 5$ bins (same approach as described in the Parameter optimization section). Finally, we visualized the difference between the binned-mean trajectories of the empirical and the shuffled data as a "streamplot", with the Python package matplotlib. The same approach was repeated with the fcANN-generated data. The similarity of the real and fcANN-generated flow analysis was quantified with Pearson's correlation coefficient (two-sided); p-values were obtained with permutation testing (1,000 permutations), by shuffling temporal order (for resting-state analyses) or condition labels (for task analyses; see below).

Evaluation: task-based dynamics

We used study 4 to evaluate the ability of the fcANN approach to capture and predict task-induced alterations in large-scale brain dynamics. First, runs 1, 3 and 7, investigating the passive experience and the down- and up-regulation of pain, respectively, were preprocessed with the same workflow used to preprocess studies 1-3 (Preprocessing and time-series extraction). Regional time series were grouped into "pain" and "rest" blocks, with a 6-second delay to adjust for the hemodynamic response time. All activation time frames were transformed to the fcANN projection plane obtained from study 1. Within-participant differences in the average location on the fcANN projection were calculated and visualized with radial plots, showing the participant-level mean trajectory on the projection plane from rest to pain (circles), as well as the group-level trajectory (arrow). The significance of the position difference and energy difference of the participant-level mean activations on the projection plane was tested with a two-sided permutation test (1,000 permutations), using the L2 norm of the two-dimensional position difference and the absolute energy difference, respectively, as test statistics, and randomly swapping the conditions within each participant.

To further highlight the difference between the task and rest conditions, a "flow analysis" was performed to investigate the dynamic trajectory differences between the conditions rest and pain. The analysis method was identical to the flow analysis of the resting state data (Evaluation: resting state dynamics). First, we calculated the direction in the projection plane between each successive TR during the rest conditions (a vector on the fcANN projection plane for each TR transition). Next, we obtained two-dimensional binned means for the x and y coordinates of these transition vectors (pooled across all participants), calculated over a two-dimensional grid of 100×100 uniformly distributed bins in the [-6,6] range (arbitrary units) and applied Gaussian smoothing with $\sigma = 5$ bins. The same procedure was repeated for the pain condition and the difference in the mean directions between the two conditions was visualized as "streamplots" (using Python's matplotlib). We used the same approach to quantify the difference in characteristic state transition trajectories between the up- and downregulation conditions. The empirically estimated trajectory differences (from real fMRI data) were contrasted to the trajectory differences predicted by the fcANN model from study 1. The similarity between real and simulated flow maps was quantified with Pearson's correlation coefficient (two-sided), and significance was assessed via permutation testing (1,000 permutations) by randomly swapping condition labels within participants.

To obtain fcANN-simulated state transitions in resting conditions, we used the stochastic relaxation procedure (Eq. (2)), with μ set to zero. To simulate the effect of pain-related activation on large-scale brain dynamics, we set μ_i during the stochastic relaxation procedure to a value representing pain-elicited activity in region i. The region-wise activations were obtained calculating the parcel-level mean activations from the meta-analytic pain activation map from (Zunhammer *et al.*, 2021), which contained Hedges' g effect sizes from an individual participant-level meta-analysis of 20 pain studies, encompassing a total of n=603 participants. The whole activation map was scaled with five different values ranging from 10^{-3} to 10^{-1}, spaced logarithmically, to investigate various signal-to-noise scenarios. We obtained the activity patterns of 10^5 iterations from this stochastic relaxation procedure and calculated the state transition trajectories with the same approach used with the empirical data. Next we calculated the fcANN-generated difference between the rest and pain conditions and compared it to the actual difference through a permutation test with 1,000 permutations, randomly swapping the conditions within each participant in the real data and using Pearson's correlation coefficient between the real (permuted) and fcANN-generated flow maps as the test statistic. From the five investigated signal-to-noise values, we chose the one that provided the highest similarity to the real pain vs. rest trajectory difference.

When comparing the simulated and real trajectory differences between pain up- and downregulation, we used the same procedure, with two differences. First, when calculating the simulated state transition vectors for the self-regulation conditions, we used the same procedure as for the pain condition, but introduced and additional signal in the nucleus accumbens, with a negative and positive sign, for up- and downregulation, respectively. We did not optimize the signal-to-noise ratio for the nucleus accumbens signal but, instead, simply used the value optimized for the pain vs. rest contrast (For a robustness analysis, see Supplementary figure 9).

Clinical data

To demonstrate the sensitivity of the fcANN approach to clinically relevant alterations of large-scale brain dynamics in autism spectrum disorder (ASD), we obtained data from n=172 individuals, acquired at the New York University Langone Medical Center (NYU) as shared in the Autism Brain Imaging Data Exchange dataset (study 7: ABIDE, (Di Martino *et al.*, 2014)). We focused on the largest ABIDE imaging center to ensure that results were not biased by center effects. We excluded high-motion cases similarly to studies 1–4, i.e., by scrubbing volumes with FD > 0.15 and excluding participants with >50% of data scrubbed. Time-series data were pooled and visualized on the fcANN projection of study 1, separately for ASD and control participants. Next, for each participant, we grouped time frames from the regional time-series data according to the corresponding attractor states (obtained with the fcANN model from study 1) and averaged time frames corresponding to the same attractor state to calculate participant-level mean attractor activations. We assessed mean attractor-activity differences between groups with a two-sided permutation test, randomly reassigning group labels 50,000 times. Reported effect sizes in the clinical tables are mean activation differences. Note that activation time series were standard-scaled independently for each region, so effect size can be interpreted as the proportion of regional variability. We adjusted the significance threshold with a Bonferroni-correction, accounting for tests across 4 states and 122 regions, resulting in $\alpha = 0.0001$. Finally, we calculated trajectory differences between the two groups, as predicted by the group-specific fcANNs (initialized with the ASD and TDC connectomes), and - similarly to the approach used in study 4 - contrasted the fcANN predictions with trajectory differences observed in real rsfMRI data. As in the previous flow analyses, we tested the significance of the similarity (Pearson's correlation) between predicted and observed trajectory differences with a two-sided permutation test (1,000 permutations), by shuffling group labels.

Methods References

Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., & Varoquaux, G. (2014). Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics*, *8*, 14.

Avants, B. B., Tustison, N. J., Song, G., Cook, P. A., Klein, A., & Gee, J. C. (2011). A reproducible evaluation of ANTs similarity metric performance in brain image registration. *Neuroimage*, *54*(3), 2033–2044.

Behzadi, Y., Restom, K., Liau, J., & Liu, T. T. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage*, *37*(1), 90–101.

Bellec, P., Rosa-Neto, P., Lyttelton, O. C., Benali, H., & Evans, A. C. (2010). Multi-level bootstrap analysis of stable clusters in resting-state fMRI. *Neuroimage*, *51*(3), 1126–1139.

Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*(3), 162–173.

Craddock, C., Sikka, S., Cheung, B., Khanuja, R., Ghosh, S. S., Yan, C., Li, Q., Lurie, D., Vogelstein, J., Burns, R., & others. (2013). Towards automated analysis of connectomes: The configurable pipeline for the analysis of connectomes (c-pac). *Front Neuroinform*, *42*(10.3389).

Dadi, K., Rahim, M., Abraham, A., Chyzhyk, D., Milham, M., Thirion, B., Varoquaux, G., Initiative, A. D. N., & others. (2019). Benchmarking functional connectome-based predictive models for resting-state fMRI. *NeuroImage*, *192*, 115–134.

Di Martino, A., Yan, C.-G., Li, Q., Denio, E., Castellanos, F. X., Alaerts, K., Anderson, J. S., Assaf, M., Bookheimer, S. Y., Dapretto, M., & others. (2014). The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular Psychiatry*, *19*(6), 659–667.

Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., Kent, J. D., Goncalves, M., DuPre, E., Snyder, M., & others. (2019). fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nature Methods*, *16*(1), 111–116.

Friston, K. J., Williams, S., Howard, R., Frackowiak, R. S., & Turner, R. (1996). Movement-related effects in fMRI time-series. *Magnetic Resonance in Medicine*, *35*(3), 346–355.

Gorgolewski, K., Burns, C. D., Madison, C., Clark, D., Halchenko, Y. O., Waskom, M. L., & Ghosh, S. S. (2011). Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python. *Frontiers in Neuroinformatics*, *5*, 13.

Jenkinson, M., Beckmann, C. F., Behrens, T. E., Woolrich, M. W., & Smith, S. M. (2012). Fsl. *Neuroimage*, *62*(2), 782–790.Ledoit, O., & Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, *88*(2), 365–411. 10.1016/s0047-259x(03)00096-4

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage*, *59*(3), 2142–2154.

Satterthwaite, T. D., Elliott, M. A., Gerraty, R. T., Ruparel, K., Loughead, J., Calkins, M. E., Eickhoff, S. B., Hakonarson, H., Gur, R. C., Gur, R. E., & others. (2013). An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *Neuroimage*, *64*, 240–256.

Spisak, T., Kincses, B., Schlitt, F., Zunhammer, M., Schmidt-Wilcke, T., Kincses, Z. T., & Bingel, U. (2020). Pain-free resting-state functional brain connectivity predicts individual pain sensitivity. *Nature Communications*, *11*(1). 10.1038/s41467-019-13785-z

Tor D., W. (2011). NeuroSynth: a new platform for large-scale automated synthesis of human functional neuroimaging data. *Frontiers in Neuroinformatics*, *5*. 10.3389/conf.fninf.2011.08.00058

Woo, C.-W., Roy, M., Buhle, J. T., & Wager, T. D. (2015). Distinct brain systems mediate the effects of nociceptive input and self-regulation on pain. *PLoS Biology*, *13*(1), e1002036.

Woo, C.-W., Roy, M., Buhle, J. T., & Wager, T. D. (2015). Distinct Brain Systems Mediate the Effects of Nociceptive Input and Self-Regulation on Pain. *PLoS Biology*, *13*(1), e1002036. 10.1371/journal.pbio.1002036

Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., & Wager, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods*, *8*(8), 665–670. 10.1038/nmeth.1635

Zunhammer, M., Spisák, T., Wager, T. D., & Bingel, U. (2021). Meta-analysis of neural systems underlying placebo analgesia from individual participant fMRI data. *Nature Communications*, *12*(1), 1391.