

Data augmentation based on dynamical systems for the classification of brain states

Yonatan Sanz Perl^{a,b,c,1}, Carla Pallavicini^{a,d,1}, Ignacio Perez Ipiña^a, Morten Kringelbach^{e,f}, Gustavo Deco^{g,h}, Helmut Laufsⁱ, Enzo Tagliazucchi^{a,b,*}

^a Department of Physics, University of Buenos Aires, Argentina

^b National Scientific and Technical Research Council (CONICET), Buenos Aires, Argentina

^c Universidad de San Andrés, Vito Dumas 284 (B1644BID), Buenos Aires, Argentina

^d Fundación para la lucha contra las enfermedades neurológicas de la infancia (FLENI), Montaneses 2325, C1428, AQK, Buenos Aires, Argentina

^e Department of Psychiatry, University of Oxford, Oxford, UK

^f Center for Music in the Brain (MIB), Dept. of Clinical Medicine, Aarhus University, Denmark

^g Center for Brain and Cognition, Computational Neuroscience Group, Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

^h Institució Catalana de la Recerca i Estudis Avançats (ICREA), Universitat Pompeu Fabra, Barcelona, Spain

ⁱ Department of Neurology, University of Kiel, Kiel, Germany

ARTICLE INFO

Article history:

Received 15 January 2020

Revised 26 March 2020

Accepted 29 June 2020

Available online 10 July 2020

Keywords:

Neuroimaging

Dynamical systems

Machine learning

Data augmentation

Brain states

ABSTRACT

The application of machine learning algorithms to neuroimaging data shows great promise for the classification of physiological and pathological brain states. However, classifiers trained on high dimensional data are prone to overfitting, especially for a low number of training samples. We describe the use of whole-brain computational models for data augmentation in brain state classification. Our low dimensional model is based on nonlinear oscillators coupled by the empirical structural connectivity of the brain. We use this model to enhance a dataset consisting of functional magnetic resonance imaging recordings acquired during all stages of the human wake-sleep cycle. After fitting the model to the average functional connectivity of each state, we show that the synthetic data generated by the model yields classification accuracies comparable to those obtained from the empirical data. We also show that models fitted to individual subjects generate surrogates with enough information to train classifiers that present significant transfer learning accuracy to the whole sample. Whole-brain computational modeling represents a useful tool to produce large synthetic datasets for data augmentation in the classification of certain brain states, with potential applications to computer-assisted diagnosis and prognosis of neuropsychiatric disorders.

© 2020 Elsevier Ltd. All rights reserved.

1. Introduction

The discovery of non-invasive neuroimaging tools opened the way to the inference of the hidden brain states that are associated with observable behaviors. For this purpose, techniques such as functional magnetic resonance imaging (fMRI) provide high dimensional spatiotemporal data that can be used as the input for machine learning classifiers [1]. In these algorithms the parameters are learned from a training sample, and the resulting accuracy is then estimated from out-of-the-sample data. A sufficiently large number of examples is critical for successful training (i.e. avoid-

ing overfitting), but the availability of neuroimaging data can be limited for certain rare neuropsychiatric conditions and for classifiers aimed at distinguishing between several groups of patients. While pooling data acquired in different laboratories can help alleviate this issue, it has been shown that heterogeneous experimental conditions can reduce the accuracy of the classifiers [2].

Data augmentation is a technique based on applying certain transformations to the available data with the objective of producing new surrogate training examples. In the case of image classification, for instance, these transformations may include rotations and shear mappings [3]. It is less obvious how to choose transformations that produce meaningful surrogate examples in the case of high dimensional spatiotemporal data, such as that provided by fMRI experiments. Faced with a similar problem, Tubaro and Mindlin recently proposed the use of low dimensional dynamical

* Corresponding author.

E-mail address: tagliazucchi.enzo@googlemail.com (E. Tagliazucchi).

¹ These authors contributed equally to this work.

systems for data augmentation in deep learning [4]. Here, we have followed the analogous procedure of developing semi-empirical models of whole-brain activity for data augmentation in the classification of temporally extended brain states.

The computational models we developed and implemented receive as input several independent sources of empirical data, and can be optimized to reproduce observables derived from fMRI recordings [5]. A low dimensional dynamical system can be assigned to each region within a brain parcellation, and inter-regional coupling can be estimated from diffusion tensor imaging (DTI) data [6]. Using the normal mode of a Hopf bifurcation results in local dynamics with a transition from a fixed point towards a limit cycle, and in global dynamics coupled by the density of long-range white matter tracts [7,8]. Finally, to reduce the dimension of the models, the local bifurcation parameters can be constrained by different functionally coherent brain systems, known as resting state networks (RSN) [9].

In the following, we show that these models reproduce the empirical correlation matrices between regional fMRI time series (also known as functional connectivity [FC] matrices), and that surrogate instances of FC matrices can be used for data augmentation in the problem of classifying the different stages of the human wake-sleep cycle. For this, synthetic time series were generated from the low dimensional models fitted to average and individual FC, which were used afterwards as input for multivariate random forest classifiers.

2. Material and methods

2.1. Participants and experimental protocol

A cohort of 63 healthy subjects participated (36 females, mean \pm SD age of 23 ± 43.3 years) with the experimental protocol approved by the local ethics committee (Goethe-Universität Frankfurt, Germany, protocol number: 305/07). Written informed consent was obtained from all participants. All experiments were conducted in accordance with the relevant guidelines and regulations, including the Declaration of Helsinki.

Within half an hour of 7 p.m. participants entered the scanner. The day of the study all participants reported a wake-up time between 5: 00 a.m. and 11: 00 a.m., and a sleep onset time between 10: 00 p.m. and 2: 00 a.m. for the night prior to the experiment. Participants were asked to relax, close their eyes and not fight the onset of sleep and, and their resting state were simultaneously measured during at least 52 m with a simultaneous combination of EEG and fMRI.

2.2. Simultaneous fMRI and EEG data collection

Electroencephalography (EEG) and electromyography (EMG) was acquired with an optimized polysomnographic setting for sleep staging. According to the rules of the American Academy of Sleep Medicine [10], the scalp potentials measured with EEG determine the classification of sleep into 4 stages (wakefulness, N1, N2 and N3 sleep). Previous publications based on this dataset can be referenced for further details [11].

2.3. Structural connectivity

Structural connectivity (SC) was obtained applying diffusion tensor imaging (DTI) to diffusion weighted imaging (DWI) recordings from 16 healthy right-handed participants (11 men and 5 women, mean age: 24.75 ± 2.54 years) recruited online at Aarhus University, Denmark. For each participant, a 90×90 SC matrix was obtained representing the density of white matter fiber tracts between regions of interest. The connectivity probability from a seed

voxel i to another voxel j was defined as the proportion of fibers passing through voxel i that reached voxel j (sampling of 5000 streamlines per voxel) [12]. All the voxels in each region of the Automated Anatomical Labeling atlas (AAL [13]) were seeded (i.e. both grey and white matter voxels were considered). The connectivity probability P_{ij} from region i to region j was computed as the number of sampled fibers in region i that connected the two regions, divided by $5000 n$, where n represents the number of voxels in region i . The resulting matrices were computed as the average across voxels within each region of interest in the AAL atlas, thresholded at 0.1 % (i.e. a minimum of five streamlines) and normalized by the number of voxels in the region. Finally, the data was averaged across all participants.

2.4. Whole-brain models

We implemented a network of nonlinear oscillators coupled by the SC. The key neurobiological assumption to implement this kind of models is that dynamics of macroscopic neural masses can range from fully synchronous to a stable asynchronous state governed by random fluctuations. Thus, each oscillator was modeled using a normal form of a Hopf bifurcation and represents the dynamics at each one of the 90 brain regions in the AAL atlas. In this type of bifurcation the qualitative nature of the solutions changes from a stable fixed point in phase space towards a limit cycle, allowing the model to represent the emergence of self-sustained oscillations. We also assume that fMRI can capture the dynamics from both regimes with sufficient fidelity to be modeled by the equations.

The local dynamics of brain region j were modeled by the complex-valued equation:

$$\frac{dz_j}{dt} = (a + i\omega_j)z_j - z_j|z_j|^2 \quad (1)$$

In this equation, z is a complex-valued variable ($z_j = x_j + iy_j$), and ω_j is the intrinsic oscillation frequency of node j . The parameter a is known as the bifurcation parameter, and the dynamical scenario changes as follows as a function of a : for $a > 0$ a limit cycle exists giving rise to self-sustained oscillations with frequency $f_j = \omega_j/2\pi$; for $a < 0$ the phase space presents a unique stable fixed point at $z_j = 0$, thus the system decays asymptotically towards this point [7]. In full form, the coupled differential equations of the model coupled by a term weighted by the SC are the following:

$$\frac{dx_j}{dt} = (a - x_j^2 - y_j^2)x_j - \omega_j y_j + G \sum C_{ij}(x_i - x_j) + \beta \eta_j \quad (2)$$

$$\frac{dy_j}{dt} = (a - x_j^2 - y_j^2)y_j + \omega_j x_j + G \sum C_{ij}(y_i - y_j) + \beta \eta_j$$

Where nodes i and j were coupled by C_{ij} (the i, j entry of the SC matrix). β_j was fixed at 0.04 and represents the scaling factor of additive gaussian noise (η_j) at each node. While the parameter G represents a global factor that scales the SC equally for all the nodes to ensure oscillatory dynamics for $a > 0$, the obtained empirical SC matrix was globally scaled to a maximum of 0.2 (weak coupling assumption). These equations were integrated to simulate empirical fMRI signals using the Euler-Maruyama algorithm with a time step of 0.1 seconds. In this model, when a is close to the bifurcation ($a \approx 0$) the additive gaussian noise gives rise to complex dynamics as the system continuously switches between both sides of the bifurcation.

2.5. Fitting the model to the empirical data

We used the group-averaged FC as the empirical observable to be fitted by the model. The fMRI signal from each region in

the AAL atlas was filtered and then z-scored. The frequency range was defined between 0.04 – 0.07 Hz since this frequency band was shown to contain more reliable and functionally relevant information compared to other frequency bands [14]. Afterwards, the FC matrix was defined as the matrix containing the Pearson correlation coefficients between the fMRI signals from all pairs of regions of interest in the AAL atlas. Fixed-effect analysis was used to obtain group-level FC matrices, meaning that the Fisher's R-to-z transform ($z = \text{atanh}(R)$) was applied to the correlation values before averaging across participants within each state of consciousness. For each brain state, participants were selected based on the presence of uninterrupted epochs of that state lasting more than 200 samples, resulting in 15 participants.

We applied the above described model (Eq. (2)) to simulate regional fMRI signals. First, we selected the global coupling factor $G = 0.5$ by an exhaustive exploration of the homogeneous parameter space (a, G) around the Hopf bifurcation ($a \approx 0$) [15]. We then incorporated an anatomical prior based on 6 major RSN [9] with the objective of constraining how different groups of nodes could contribute to the local bifurcation parameters. In this way, we embedded the dynamics of the 90 independent regions into a 6-dimensional parameter space. Each bifurcation parameter was constructed as the linear combination of the 6 parameters associated with the RSN. Note that regions could belong to more than one RSN, and thus the bifurcation parameters could receive independent contributions from multiple RSN. We simulated 200 time samples for each subject, and then repeated the procedure described above to compute the simulated average FC matrices for each state. The goodness of fit (GoF) was determined by the structure similarity index (SSIM) [16], an image similarity metric that factors both the similarity between the image means and between their covariance structures. $SSIM = 1$ occurs when the comparison is between two identical images and $SSIM = 0$ occurs when the two images are completely different. The optimization procedure was based on genetic algorithms applied to infer the 6 parameters that maximize the goodness of fit. Further details can be found in previous work implementing the same model [17].

2.6. Multivariate machine learning classifiers and data augmented by the model

We trained random forest classifiers [18] to distinguish sleep from wakefulness based on FC matrices, using a five-fold cross-validation procedure to estimate the accuracy. Classifiers were trained to distinguish between wakefulness and a certain sleep stage, and their accuracy was then tested in the classification between wakefulness and the same as well as other sleep stages (i.e. transfer learning accuracy).

Random forest classifiers were implemented using scikit-learn (<https://scikit-learn.org/>) [19]. Briefly, the random forest algorithm builds upon the concept of a decision tree classifier, where samples are iteratively split into two branches depending on the values of their features. For each feature, a threshold is introduced so that the samples are separated in a way that maximizes a metric of the homogeneity of the class labels assigned to each branch. The algorithm stops whenever a split results in a branch where all the samples belong to the same class, or when all features were already used for a split. Since this procedure is prone to overfitting, the random forest algorithm trains an ensemble of decision trees based on a randomly chosen subset of the features, and then computes the label prediction as the majority vote across all the individual trees.

We trained random forest classifiers with 1000 decision trees and a random subset of features of size equal to the (rounded) square root of the total number of features. The quality of each split in the decision trees was measured using Gini impurity, and

the individual trees were expanded until all leaves were pure (i.e. no maximum depth). No minimum impurity decrease was enforced at each split, and no minimum number of samples was required at the leaf nodes of the decision trees (the classifier hyperparameters can be found in <https://scikit-learn.org/>).

To assess the statistical significance of the accuracy values, we trained and evaluated a total of 1000 random forest classifiers using the same features (i.e. FC matrices) but scrambling the class labels. We then constructed an empirical p-value by counting how many times the accuracy of the classifier with scrambled class labels was greater than that of the original classifier. All accuracies were determined as the area under the receiver operating characteristic curve (AUC). Subsequently, the generalizability of the classifiers to distinguish other sleep states from wakefulness was evaluated by applying both the original and scrambled classifiers, and constructing a p-value in a similar way.

We repeated the aforementioned procedure using data augmentation, given by the output of the whole-brain computational model. We optimized model parameters using the average FC matrices as the targets, and used these parameters to simulate 100 surrogate samples. The inclusion of additive noise in the model gives rise to different simulated time series for each independent run, and consequently to different FC matrices. Thus, these samples are not independent from the original data, since they are created by simulations with the same optimized parameters used to fit the real data. Nevertheless, they are transformed in a meaningful way that incorporates our knowledge in terms of the mechanistic explanation proportioned the model, i.e., in each brain state some regions seemingly operate at the border between stability and instability ($a \approx 0$), others behave as a noisy fixed point ($a < 0$), and regions are connected by the structural connectivity. Thus, this procedure guarantees that the generated surrogate data based on optimized parameters to fit one particular sleep state will represent informative and useful variations of this state.

In this way, we obtained 100 synthetic samples for each stage. Based on these surrogate samples, we trained classifiers to distinguish wakefulness from sleep and measured the accuracy of these classifiers using the empirical data. We also determined the transfer learning accuracy by evaluating the performance of the classifiers trained with surrogate data of a certain sleep stage in the problem of classifying empirical data corresponding to another sleep stage. Finally, we randomly selected three subjects from the empirical dataset and repeated this procedure using single subject FC matrices as optimization targets for the whole-brain model.

3. Results

The procedure we followed is outlined in Fig. 1. First, we combined three different sources of empirical data to inform the computational model based on coupled Hopf bifurcations (Eq. (2)). Then, we trained and evaluated random forest classifiers to distinguish different pairs of sleep stages based on empirical and synthetic fMRI data, as well as on single subject synthetic data.

The first row of Fig. 2 shows the average empirical FC matrices corresponding to wakefulness and the three stages of NREM sleep (N1, N2, N3). In these matrices, rows and columns correspond to one of the 90 regions in the AAL atlas, and the correlation coefficient between fMRI time series is indicated by the color scale. The remaining rows show the FC matrices computed for three randomly chosen subjects.

Fig. 3 contains the same information computed from the synthetic fMRI data. The main difference between the empirical and simulated matrices appeared in the contradiagonal, which corresponds to interhemispheric (or homotopic) connections (i.e. connections between two regions symmetrically located with respect to the midline), a difference consistent with the observation that

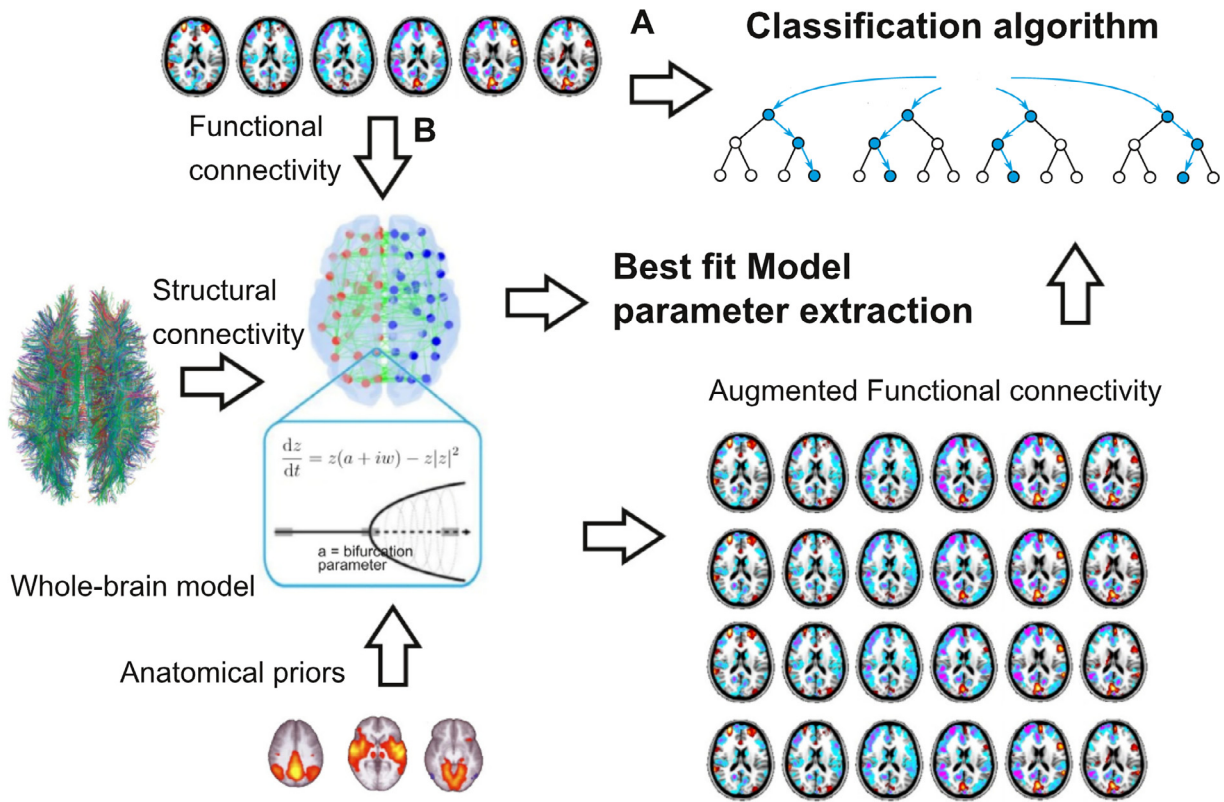


Fig. 1. Outline of the procedure followed to generate the synthetic data and train the machine learning classifiers. Three sources of empirical data informed the computational model based on coupled nonlinear oscillators (Eq. (2)). SC represented the coupling strength between oscillators, FC was used as the target function for parameter optimization with genetic algorithms, and 6 RSN determined the anatomical priors constraining local contributions to bifurcation parameters. The empirical FC was used both as input to the random forest classifiers (A) and as target function (B) in the optimization procedure. After this step the model generated surrogate samples to augment the training data.

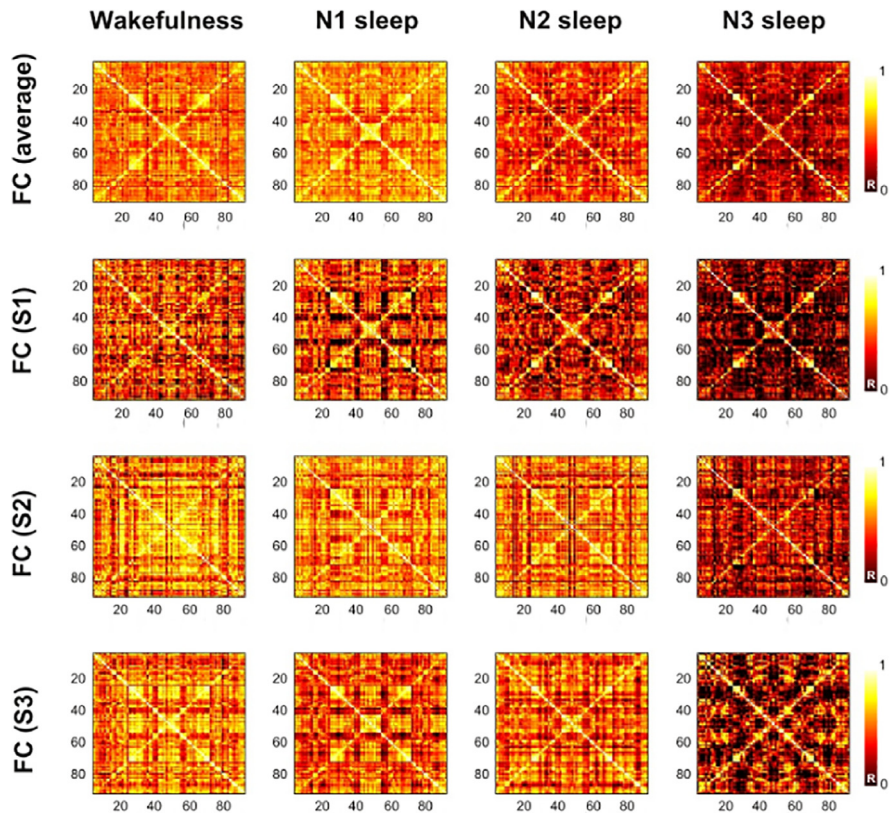


Fig. 2. Empirical FC matrices containing the correlation coefficients (R) between fMRI time series from all pairs of regions in the AAL atlas. The first row displays the average FC matrices for wakefulness, N1, N2 and N3 sleep. The other rows contain the same information for three randomly chosen individuals.

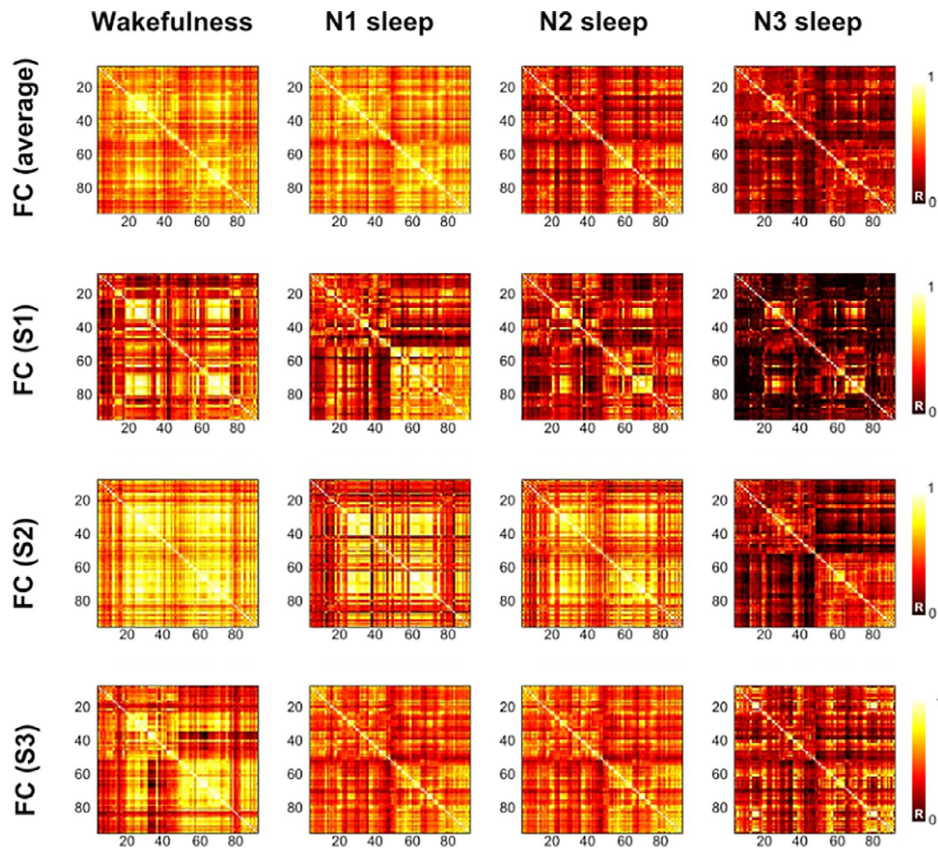


Fig. 3. FC matrices containing the correlation coefficients (R) between synthetic fMRI time series from all pairs of regions in the AAL atlas. The first row displays the optimal simulated FC matrices for wakefulness, N1, N2 and N3 sleep. The other rows contain the same information for three randomly chosen individuals.

Table 1

Goodness of fit (SSIM) between empirical and simulated FC matrices, both for the average data and for the three individual subjects (S1, S2, S3).

	W	N1	N2	N3
Average	0.40	0.40	0.41	0.37
S1	0.27	0.28	0.24	0.21
S2	0.37	0.30	0.54	0.22
S3	0.42	0.30	0.39	0.27

DTI tends to underestimate long-range fiber tracts [20]. In general, the model fit was better for the average FC compared to individual subjects (see Table 1).

Panel A of Fig. 4 shows the histograms of AUC values representing the transfer learning accuracy for random forest classifiers trained using 100 synthetic wakefulness samples and 100 synthetic N1/N2/N3 sleep samples, and evaluated in the empirical data. Each column and row indicates the sleep stage used for training and testing, respectively. For instance, the second plot of the first row contains the AUC histograms obtained in the classification between the empirical FC matrices from wakefulness and N2 sleep ($N = 15$ subjects), using the classifier trained to distinguish wakefulness from N1 sleep based on synthetic fMRI data ($N = 100$ surrogates). The histograms in red correspond to the AUC values obtained using the real data, while the histograms in blue indicate the AUC values obtained after shuffling the data labels. Label shuffled is used as a null model to obtain the p-values shown in the insets.

The matrices in Panel B of Fig. 4 summarize the average AUC obtained for all training-generalization pairs. It is clear from observing these matrices that machine learning classifiers presented

the highest transfer learning accuracy when generalizing between N2 and N3 sleep. This result was obtained using both synthetic (left) and empirical (right) data for training. The scatter plot compares the entries of both matrices, showing a positive correlation which supports the similarity between the empirical and simulated AUC matrices.

We also explored other combinations of training and testing set, including synthetic data generated with and without the ad hoc inclusion of the contradiagonal (i.e. homotopic structural connectivity), and a combination of empirical and synthetic data in different proportions. These results are shown in Fig. S1 and Fig. S2 of the Supplementary Information, for the cross validation and transfer learning, respectively.

Fig. 5 shows that data augmentation based on single subject FC matrices can be used to train machine learning classifiers that present significant accuracy in the classification of wakefulness from N2 and N3 sleep. The rows correspond to random forest classifiers trained using data generated by computational models fitted to the empirical FC of S1, S2 and S3 ($N = 100$ surrogates). The columns indicate the sleep stage to be distinguish from wakefulness. All histograms contain AUC values obtained from the evaluation of these models on the empirical data ($N = 15$ subjects), both with unshuffled (red) and shuffled (blue) class labels. The resulting p-values indicate that classifiers trained using synthetic data from individual subjects can successfully generalize to the whole sample in the classification of wakefulness vs. N2 and N3 sleep, but not vs. N1 sleep.

4. Discussion

One of the main limitations for the training of machine learning classifiers is the amount of available data. Training is gener-

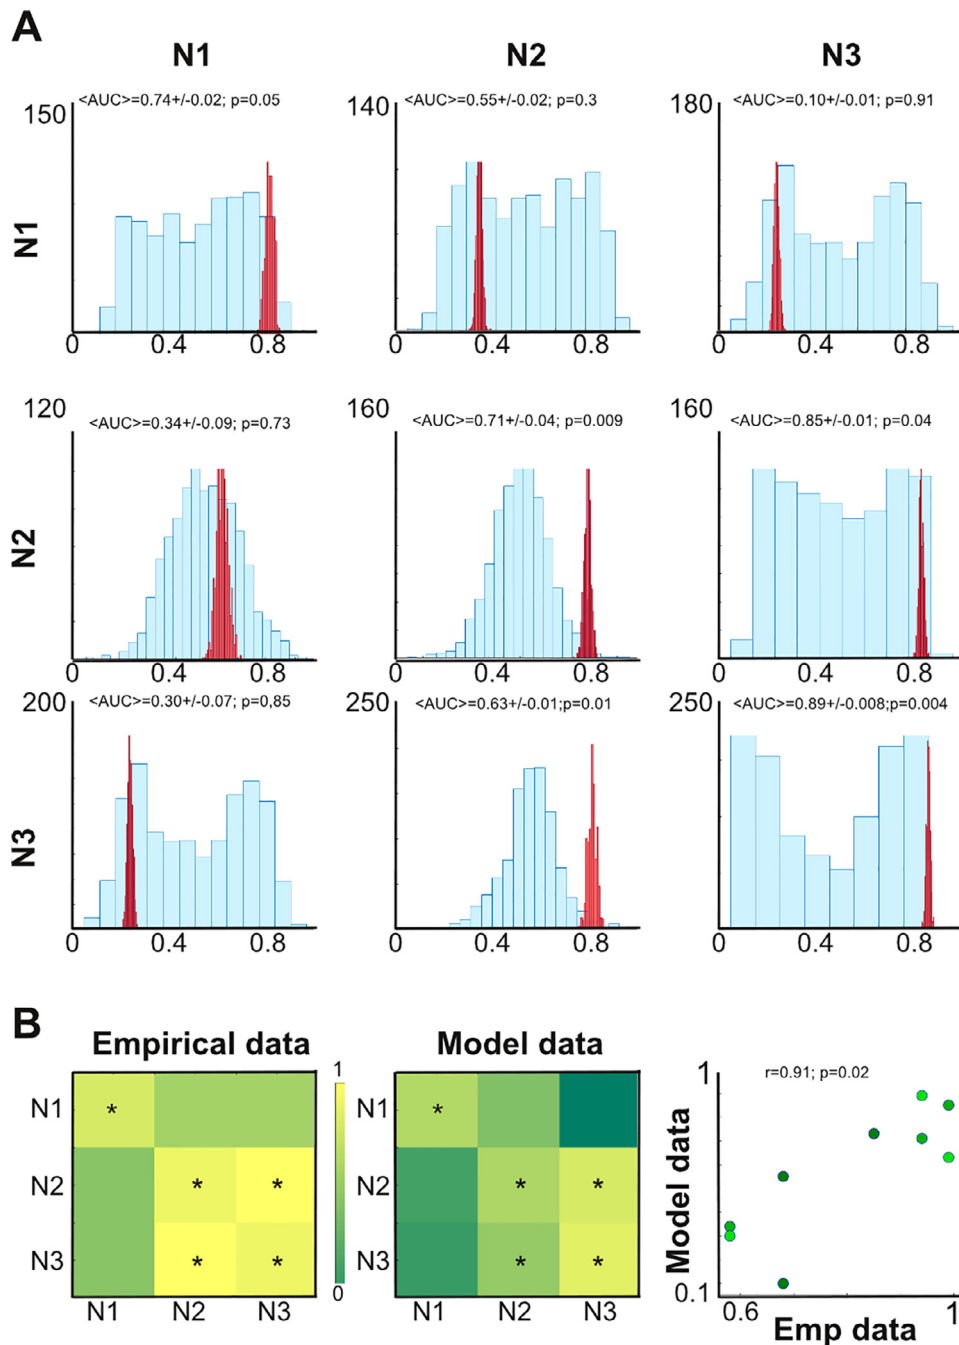


Fig. 4. A) Histograms of AUC values for the random forest classifiers trained to distinguish wakefulness vs. the sleep stages indicated in the rows, and tested in the classification of wakefulness vs. the sleep stages indicated in the columns. All classifiers were trained using synthetic FC matrices fitted to the average FC matrices ($N = 100$ surrogates) and evaluated using the empirical data ($N = 15$). B) Matrices containing the average AUC values obtained for the random forest classifiers trained using the empirical (left) and the synthetic (right) data. Asterisks indicate statistically significant AUC values (determined by the label shuffling procedure). The scatter plot contains the entries of the “empirical data” vs. the “model data” matrices.

ally successful provided sufficient data and informative features, but overfitting can drastically reduce generalization performance if only few training samples are available. Data augmentation techniques can attenuate this problem by introducing certain transformations (such as shear mappings rotations, in the case of images); however, it is not clear how complex spatiotemporal data should be transformed to create meaningful surrogate samples. Another important application of data augmentation consists of enriching datasets with a very small number of samples with surrogates presenting sufficient variability to allow the training of machine learning classifiers. Following recently published work [4], we showed

that low dimensional dynamical systems fitted to empirical observables can be successfully applied for data augmentation with the purpose of brain state classification. We note that surrogate time series can be produced by methods not based on dynamical systems (e.g. [21]). However, models such as the one we employed can represent advantages in terms of interpretation and conceptual clarity. They can also be tailored to train classifiers using synthetic samples deviating in useful ways from the available experimental data. Finally, their semi-empirical nature facilitates the transition towards the single subject level. In the following, we discuss these advantages in the context of the present work.

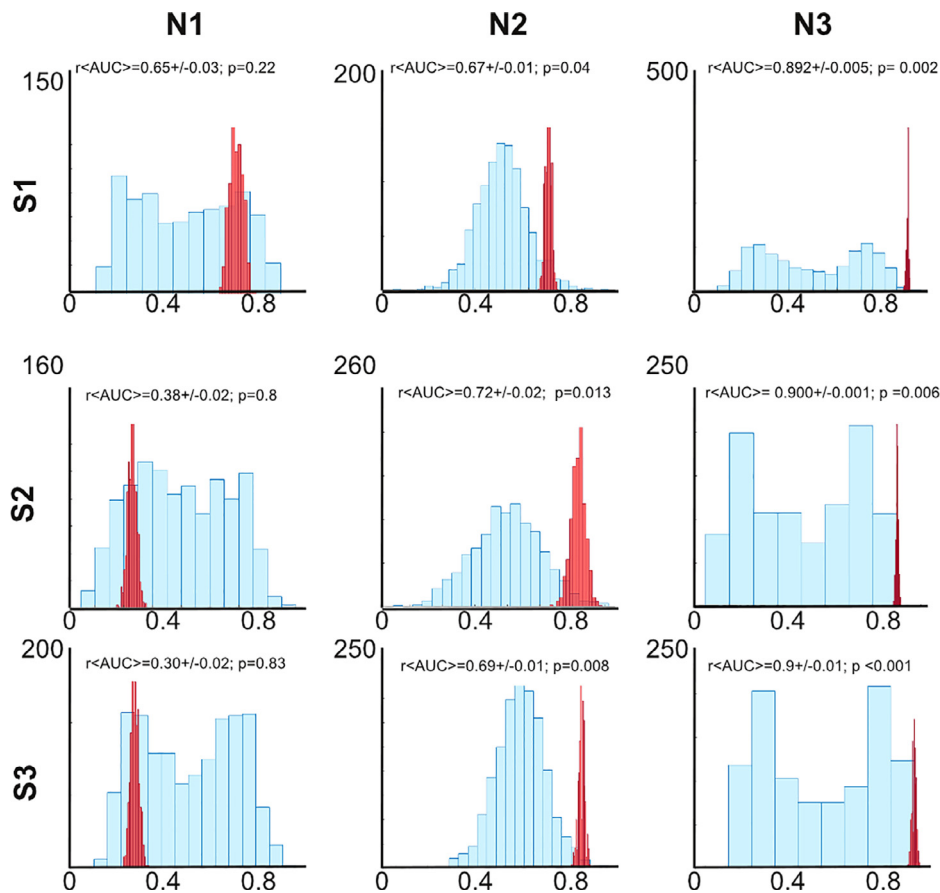


Fig. 5. Histograms of AUC values for random forests classifiers trained to distinguish N1, N2 and N3 from wakefulness. Classifiers were trained using synthetic FC matrices based on data from individual subjects ($N = 100$ surrogates) and evaluated on empirical samples ($N = 15$ subjects). Histograms in red correspond to data without label shuffling, while blue indicates AUC after label shuffling. Insets contain the mean AUC \pm SD and the associated p-values.

The computational models we used to generate synthetic data did not strive for biological realism, instead, we decided to focus on the simplest dynamics that could present the kind of behavior needed to provide the classifiers with representative surrogates for training. Based on this data, the classifiers were capable of inferring the optimal separation boundary between classes and presented significant transfer accuracy for the generalization between N2 and N3, which was expected considering the high behavioral and physiological similarities between these stages compared to N1 sleep [10]. The classifier transfer learning accuracy matrices obtained from synthetic and empirical data were very similar, supporting the conceptual validity of our simple model for whole-brain dynamics, which could be expected from previous work based on similar dynamics [7,8,17]. Low complexity models can simultaneously preserve the informativeness of the surrogates while allowing the exploration of a small set of interpretable parameters.

The classification of brain states based on fMRI recordings is a promising tool for the automated diagnosis and prognosis of certain neuropsychiatric patients [22], however, this promise is frequently undermined by small sample sizes [23]. Building databases of fMRI recordings can be costly and time consuming for diseases that are rare or difficult to investigate with neuroimaging. Also, developing algorithms for differential diagnosis requires multilabel classifiers, which further reduces the number of samples per class. A possible solution to this issue is gathering data from multiple research groups; however, different scanners and imaging sequences can be critical confounds for machine learning classifiers [2]. As an alternative, we proposed that adequate data augmentation tech-

niques based on computational modeling can contribute to overcoming these limitations. We note that these are not mutually exclusive solutions, for instance, models could be used to explore parametrically how classifiers are confounded by factors related to variability in the experimental conditions.

The outcome of our model depends upon a relatively low number of parameters, which could be explored to train classifiers with surrogate samples including perturbations that represent the hypothesized outcome of certain interventions. For instance, the outcome of surgical brain resection in certain forms of epilepsy could be modeled by localized SC changes [24]. By artificially inducing these changes in the model parameters (including the structural coupling between nodes) it could be possible to produce synthetic data useful to train classifiers that can be applied to estimate the likelihood of success after the intervention. The same logic could be applied to other kinds of treatments, such as pharmacological interventions and non-invasive brain stimulation protocols, as well as to train machine learning classifier with data that simulates specific lesions, such as those arising from stroke and traumatic brain injury.

We have shown that data augmentation using models fitted to single subject FC matrices also allowed the classifiers to distinguish between wakefulness and sleep. As such, our results represent an encouraging proof of concept, but care should be exercised when attempting to generalize this result to other brain states. Since sleep is a physiological process and our population consisted of healthy participants, we expected that individual subjects could provide enough information to develop classifiers accurate at the group level. However, this cannot be taken for granted

in surrogates obtained from models fitted to individual patients, where higher inter-subject variability may arise from abnormalities in brain structure and function. Since these limitations could be informative of such abnormalities, low dimensional whole-brain models should be further explored in the context of reproducing single subject FC from the individual SC of the patients [25].

In conclusion, we have shown that dynamical systems constitute a valuable tool for generating synthetic spatiotemporal data based on a small number of examples, a tool that can be naturally applied for data augmentation when training automated classifiers using fMRI data. Future work should study the possibility of overcoming data scarcity in other systems that can be modeled by simple dynamics, contributing to the fruitful cross-fertilization of artificial intelligence and physics.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Yonatan Sanz Perl: Conceptualization, Investigation, Methodology, Visualization, Writing - original draft. **Carla Pallavicini:** Investigation, Methodology, Visualization. **Ignacio Perez Ipiña:** Investigation, Methodology, Visualization. **Morten Kringelbach:** Methodology, Writing - review & editing. **Gustavo Deco:** Methodology, Writing - review & editing. **Helmut Laufs:** Data curation, Funding acquisition, Project administration, Writing - review & editing. **Enzo Tagliazucchi:** Conceptualization, Project administration, Supervision, Visualization, Writing - original draft.

Acknowledgments

Authors acknowledge funding from Agencia Nacional De Promocion Cientifica Y Tecnologica (Argentina), grant PICT-2018-03103.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.chaos.2020.110069](https://doi.org/10.1016/j.chaos.2020.110069).

References

- [1] Richiardi J, Eryilmaz H, Schwartz S, Vuilleumier P, Van De Ville D. Decoding brain states from fMRI connectivity graphs. *Neuroimage* 2011;56(2):616–26.
- [2] Kassraian-Fard P, Matthias C, Balsters JH, Maathuis MH, Wenderoth N. Promises, pitfalls, and basic guidelines for applying machine learning classifiers to psychiatric imaging data, with autism as an example. *Front Psychiatry* 2016;7:177.
- [3] Perez L, Wang J. The effectiveness of data augmentation in image classification using deep learning. *arXiv:1712.04621* 2017.
- [4] Tubaro P, Mindlin G. A dynamical system as the source of augmentation in a deep learning problem. *Chaos Soliton Fract* 2019;2:100012.
- [5] Breakspear M. Dynamic models of large-scale brain activity. *Nat Neurosci* 2017;20(3):340.
- [6] Hagmann P, Cammoun L, Gigandet X, Meuli R, Honey CJ, Wedeen VJ, et al. Mapping the structural core of human cerebral cortex. *PLoS Biol* 2008;6(7):e159.
- [7] Deco G, Kringelbach ML, Jirsa VK, Ritter P. The dynamics of resting fluctuations in the brain: metastability and its dynamical cortical core. *Sci Rep* 2017;7(1):3095.
- [8] Deco G, Cabral J, Saenger VM, Boly M, Tagliazucchi E, Laufs H, et al. Perturbation of whole-brain dynamics in silico reveals mechanistic differences between brain states. *NeuroImage* 2018;169:46–56.
- [9] Damoiseaux JS, Rombouts S, Barkhof F, Scheltens P, Stam CJ, Smith SM, et al. Consistent resting-state networks across healthy subjects. *Proc Natl Acad Sci* 2006;103(37):13848–53.
- [10] Berry RB, Brooks R, Gamaldo CE, Harding SM, Marcus C, Vaughn BV, et al. The aasm manual for the scoring of sleep and associated events: Rules, Terminology and Technical Specifications, Darien, Illinois. *Am Acad Sleep Med* 2012;176.
- [11] Tagliazucchi E, Laufs H. Decoding wakefulness levels from typical fMRI resting-state data reveals reliable drifts between wakefulness and sleep. *Neuron* 2014;82(3):695–708.
- [12] Behrens TE, Berg HJ, Jbabdi S, Rushworth MF, Woolrich MW. Probabilistic diffusion tractography with multiple fibre orientations: what can we gain? *Neuroimage* 2007;34(1):144–55.
- [13] Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, et al. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage* 2002;15(1):273–89.
- [14] Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, et al. Frequencies contributing to functional connectivity in the cerebral cortex in “resting-state” data. *Am J Neuroradiol* 2001;22(7):1326–33.
- [15] Jobst BM, Hindriks R, Laufs H, Tagliazucchi E, Hahn G, Ponce-Alvarez A, et al. Increased stability and breakdown of brain effective connectivity during slow-wave sleep: mechanistic insights from whole-brain computational modelling. *Sci Rep* 2017;7(1):1–16.
- [16] Wang Z, Bovik AC, Sheikh HR, Simoncelli EP, et al. Image quality assessment: from error visibility to structural similarity. *IEEE Trans Image Process* 2004;13(4):600–12.
- [17] Ipiña IP, Perl YS, Kringelbach ML, Kehoe PD, Deco G, Laufs H, et al. Modeling the relationship between regional activation and functional connectivity during wakefulness and sleep. *NeuroImage*, 2020, p. 116833.
- [18] Breiman L. Random forests. *Mach Learn* 2001;45(1):5–32.
- [19] Abraham A, Pedregosa F, Eickenberg M, Gervais P, Mueller A, Kossaifi J, et al. Machine learning for neuroimaging with scikit-learn. *Front Neuroinform* 2014;8:14.
- [20] Reveley C, Seth AK, Pierpaoli C, Silva AC, Yu D, Saunders RC, et al. Superficial white matter fiber systems impede detection of long-range cortical connections in diffusion mr tractography. *Proc Natl Acad Sci* 2015;112(21):E2820–8.
- [21] Laumann TO, Snyder AZ, Mitra A, Gordon EM, Gratton C, Adeyemo B, et al. On the stability of bold fMRI correlations. *Cereb Cortex* 2016;27(10):4719–32.
- [22] Klöppel S, Abdulkadir A, Jack Jr CR, Koutsouleris N, Mourão-Miranda J, Vemuri P. Diagnostic neuroimaging across diseases. *Neuroimage* 2012;61(2):457–63.
- [23] Snack H, Kahn R. Detecting neuroimaging biomarkers for psychiatric disorders: sample size matters. *Front Psychiatry* 2016;7:50.
- [24] Taylor PN, Kaiser M, Dauwels J. Structural connectivity based whole brain modelling in epilepsy. *J Neurosci Methods* 2014;236:51–7.
- [25] Falcon MI, Jirsa V, Solodkin A. A new neuroinformatics approach to personalized medicine in neurology: the virtual brain. *Curr Opin Neurol* 2016;29(4):429.