



HAL
open science

Functional Magnetic Resonance Imaging data augmentation through conditional ICA

Badr Tajini, Hugo Richard, Bertrand Thirion

► **To cite this version:**

Badr Tajini, Hugo Richard, Bertrand Thirion. Functional Magnetic Resonance Imaging data augmentation through conditional ICA. MICCAI 2021 - 24th International Conference on Medical Image Computing and Computer Assisted Intervention, Sep 2021, Strasbourg, France. hal-03284313

HAL Id: hal-03284313



<https://hal.archives-ouvertes.fr/hal-03284313>

Submitted on 12 Jul 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Functional Magnetic Resonance Imaging data augmentation through conditional ICA

Badr Tajini *, Hugo Richard *, Bertrand Thirion 

Inria, CEA, Université Paris-Saclay, France

Abstract. Advances in computational cognitive neuroimaging research are related to the availability of large amounts of labeled brain imaging data, but such data are scarce and expensive to generate. While powerful data generation mechanisms, such as Generative Adversarial Networks (GANs), have been designed in the last decade for computer vision, such improvements have not yet carried over to brain imaging. A likely reason is that GANs training is ill-suited to the noisy, high-dimensional and small-sample data available in functional neuroimaging. In this paper, we introduce Conditional Independent Components Analysis (Conditional ICA): a fast functional Magnetic Resonance Imaging (fMRI) data augmentation technique, that leverages abundant resting-state data to create images by sampling from an ICA decomposition. We then propose a mechanism to condition the generator on classes observed with few samples. We first show that the generative mechanism is successful at synthesizing data indistinguishable from observations, and that it yields gains in classification accuracy in brain decoding problems. In particular it outperforms GANs while being much easier to optimize and interpret. Lastly, Conditional ICA enhances classification accuracy in eight datasets without further parameters tuning.

Keywords: Conditional ICA · Data generation · Decoding studies.

1 Introduction

As a non-invasive brain imaging technique, task fMRI records brain activity while participants are performing specific cognitive tasks. Univariate statistical methods, such as general linear models (GLMs) [7] have been successfully applied to identifying the brain regions involved in specific tasks. However such methods do not capture well correlations and interactions between brain-wide measurements. By contrast, classifiers trained to *decode* brain maps, i.e to discriminate between specific stimulus or task types [15, 25, 27], take these correlations into account. The same framework is also popular for individual imaging-based diagnosis.

However, the large sample-complexity of these classifiers currently limits their accuracy. To tackle this problem, data generation is an attractive approach, as

* These authors contributed equally to this work

it could potentially compensate for the shortage of data. Generative Adversarial Networks (GANs) are promising generative models [8]. However, GANs are ill-suited to the noisy, high-dimensional and small-sample data available in functional neuroimaging. Furthermore the training of GANs is notoriously unstable and there are many hyper-parameters to tune.

In this work, we introduce Conditional ICA: a novel data augmentation technique using ICA together with conditioning mechanisms to generate surrogate brain imaging data and improve image classification performance. Conditional ICA starts from a generative model of resting state data (unconditional model), that is fine-tuned into a conditional model that can generate task data. This way, the generative model for task data benefits from the abundant resting state data and can be trained with few labeled samples. We first show that the generative model of resting state data shipped in Conditional ICA produces samples that neither linear nor non-linear classifiers are able to distinguish. Then we benchmark Conditional ICA as a generative model of task data against various augmentation methods including GANs and conditional GANs on their ability to improve classification accuracy on a large task fMRI dataset. We find that Conditional ICA yields highest accuracy improvements. Lastly, we show on 8 different datasets that the use of Conditional ICA results in systematic improvements in classification accuracy ranging from 1% to 5%.

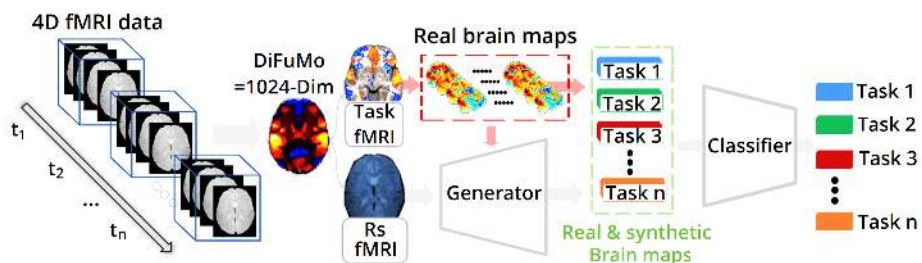


Fig. 1: **Conditional ICA approach.** Our method aims to generate surrogate data from Task and Rest fMRI data by synthesizing statistical maps that qualitatively fit the distribution of the original maps. These can be used to improve the accuracy of machine learning models that identify contrasts from the corresponding brain activity patterns.

2 Methods

Notations We write matrices as bold capital letters, vectors as small bold letters. \mathbf{X}^\dagger refers to the Moore-Penrose pseudo inverse of matrix \mathbf{X} , $\text{tr}(\mathbf{X})$ refers to the trace of matrix \mathbf{X} and \mathbf{I}_k refers to the identity matrix in $\mathbb{R}^{k,k}$. $\mathbf{0}_k$ refers to the null vector in \mathbb{R}^k .

Spatial Dimension reduction. The outline of the proposed approach is presented in Fig.1. While brain maps are high-dimensional, they span a smaller space than that of the voxel grid. For the sake of tractability, we reduce the dimension of the data by projecting the voxel values on the high-resolution version of the Dictionaries of Functional Modes *DiFuMo* atlas [3], i.e. with $p = 1024$ components. The choice of dimension reduction technique generally has an impact on the results. However we consider this question to be out of the scope of the current study and leave this to future work.

Unconditional generative models (resting state data generation) Given a large-scale resting-state dataset \mathbf{X}^{rest} in $\mathbb{R}^{p,n}$ where n is the number of images (samples) and $p = 1024$ the number of components in the atlas, let us consider how to learn its distribution. Assuming a Gaussian distribution is standard in this setting, yet, as shown later, it misses key distributional features. Moreover, we consider a model that subsumes the distribution of any type of fMRI data (task or rest): a linear mixture of $k \leq p$ independent temporal signals. We therefore use temporal ICA to learn a dimension reduction and unmixing matrix $\mathbf{W}^{rest} \in \mathbb{R}^{k,p}$ such that the k sources i.e the k components of $\mathbf{S}^{rest} = \mathbf{W}^{rest}\mathbf{X}^{rest}$ are as independent as possible.

A straightforward method to generate new rest data would be to independently sample them from the distribution of the sources. This is easy because such distribution has supposedly independent marginals. We apply an invertible quantile transform q^{rest} to the sources \mathbf{S}^{rest} so that the distribution of $\mathbf{z}^{rest} = q^{rest}(\mathbf{s}^{rest})$ has standardized Gaussian marginals. Since the distribution of \mathbf{z}^{rest} has independent marginals, it is given by $\mathcal{N}(\mathbf{0}_k, \mathbf{I}_k)$ from which we can easily sample. As shown later, this approach fails: such samples are still separable from actual rest data.

We hypothesize that this is because independence does not hold, and thus a latent structure among the marginals of the source distribution has to be taken into account. Therefore we assume that the distribution of \mathbf{z}^{rest} is given by $\mathcal{N}(\mathbf{0}_k, \mathbf{A}^{rest})$ where \mathbf{A}^{rest} is a definite positive matrix. \mathbf{A}^{rest} can easily be learned from a standard shrunk covariance estimator: $\mathbf{A}^{rest} = \mathbf{\Sigma}^{rest}(1-\alpha) + \alpha \text{tr}(\mathbf{\Sigma}^{rest})\mathbf{I}_k$ where α is given by the Ledoit-Wolf formula [14] and $\mathbf{\Sigma}^{rest}$ is the empirical covariance of \mathbf{Z}^{rest} .

Our encoding model for rest data is therefore given by $\mathbf{Z}^{rest} = q^{rest}(\mathbf{W}^{rest}\mathbf{X}^{rest})$ and we assume that the distribution of \mathbf{Z}^{rest} is $\mathcal{N}(\mathbf{0}_k, \mathbf{A}_k)$. The generative model is given by the pseudo inverse of the encoding model: $\tilde{\mathbf{X}}^{rest} = (\mathbf{W}^{rest})^\dagger (q^{rest})^{-1}(\boldsymbol{\varepsilon})$ where $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}_k, \mathbf{A}^{rest})$.


Conditional generative models (generative model for task data) While resting state datasets have a large number of samples ($10^4 \sim 10^5$), task datasets have a small number of samples ($10 \sim 10^2$). As a result, there are too few samples to learn high quality unmixing matrices. Therefore, using the unmixing matrix \mathbf{W}^{rest} learned from the resting state data, we rely on the following nonlinear generative model for brain maps in a certain class c :

$$\mathbf{x}_c = (\mathbf{W}^{rest})^\dagger q^{-1}(\boldsymbol{\varepsilon}) \quad (1)$$

with $\varepsilon \sim \mathcal{N}(\boldsymbol{\mu}_c, \boldsymbol{\Lambda})$.

In order to maximize the number of samples used to learn the parameters of the model, we assume that the quantile transform q and the latent covariance $\boldsymbol{\Lambda}$ do not depend on the class c . However, the mean $\boldsymbol{\mu}_c$, that can be learned efficiently using just a few tens of samples, depends on class c . An overview of our generative method is shown in Fig. 2.

Parameters estimation



Data generation



Performance evaluation




Fig. 2: **Conditional ICA approach in depth.** The approach proceeds by learning a temporal ICA of rest data $\mathbf{X}^{rest} \in \mathbb{R}^{p,n}$, resulting in independent sources and unmixing matrix $\mathbf{W}^{rest} \in \mathbb{R}^{k,p}$. Applying the unmixing matrix to the task data, we obtain samples in the source space $\mathbf{S}^{task} \in \mathbb{R}^{k,n}$. Afterwards, we map \mathbf{S}^{task} to a normal distribution, yielding $\mathbf{Z}^{task} \in \mathbb{R}^{k,n}$. Then, we estimate the covariance $\boldsymbol{\Lambda} \in \mathbb{R}^{k,k}$ (all classes are assumed to have the same covariance) and the class-specific means $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_C \in \mathbb{R}^k$ according to Ledoit-Wolf’s method. For each class c , we can draw random samples $\tilde{\mathbf{Z}}_c^{task} \in \mathbb{R}^{k,n_{fakes}}$ from the resulting multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_c, \boldsymbol{\Lambda})$ and obtain fake data $\tilde{\mathbf{X}}_c \in \mathbb{R}^{p,n_{fakes}}$ by applying the inverse quantile transform and re-mixing the data using the pseudo inverse of the unmixing matrix. We append these synthetic data to the actual data to create our new augmented dataset on which we train classifiers.

3 Related work

In image processing, data augmentation is part of standard toolboxes and typically includes operations like cropping, rotation, translation. On fMRI data these methods do not make much sense as brain images are not invariant to such transformations. More advanced techniques [29] are based on generative models such as GANs or variational auto-encoders [12]. Although GAN-based method are

powerful they are slow and difficult to train [2]. In appendix Table 7, we show that Conditional ICA is several orders of magnitude faster than GAN-based methods.

Our method is not an adversarial procedure. However it relates to other powerful generative models such as variational auto-encoders [12] with which it shares strong similarities. Indeed the analog of the encoding function in the variational auto-encoder is given by $e(\mathbf{x}) = \mathbf{A}^{-\frac{1}{2}}q(\mathbf{W}^{rest}\mathbf{x})$ in our model and the analog to the decoding function in the variational auto-encoder is given by $d(\mathbf{z}) = (\mathbf{W}^{rest})^\dagger q^{-1}(\mathbf{A}^{\frac{1}{2}}\mathbf{z})$ in our model. As in the variational auto-encoder, e approximately maps the distribution of the data to a standardized Gaussian distribution, while the reconstruction error defined by the difference in l2 norm $\|d(e(\mathbf{x})) - \mathbf{x}\|_2^2$ must remain small.

Another classical generative model related to ours is given by normalizing flows [13]. When \mathbf{W}^{rest} is squared (no dimension reduction in ICA), the decoding operator d is invertible (its inverse is e) making our model an instance of normalizing flows. A great property is thus the simplicity and reduced cost of data generation.

Software tools We use the FastICA algorithm [11] to perform independent component analysis and use $k = 900$ components. We use Nilearn [1] for fMRI processing, Scikit-learn [19], Numpy [10] and Scipy [28] for data manipulation and machine learning tools.

Code availability Our code is available on the following git repository: <https://github.com/BTajini/augfmri>.

4 Experiments

4.1 Dataset, data augmentation baselines and classifiers used

The unmixing matrices are learned on the rest HCP dataset [26] using 200 subjects. These data were used after standard preprocessing, including linear detrending, band-pass filtering ($[0.01, 0.1]Hz$) and standardization of the time courses. The other 8 datasets [18, 20, 20–24, 26] are obtained from the Neurovault repository [9]. The classes used in each dataset correspond to the activation maps related to the contrasts (such as “face vs tools”) present in the set of tasks of each dataset. Details are available in appendix Table 3.

We consider 5 alternative augmentation methods: *ICA*, *Covariance*, *ICA + Covariance*, *GANs* and *CGANs*. When no augmentation method is applied we use the label *Original*.

The *ICA* covariance method applies ICA to \mathbf{X}^{task} to generate unmixing matrices \mathbf{W}^{task} and sources $\mathbf{S}^{task} = \mathbf{W}^{task}\mathbf{X}^{task}$. To generate a sample $\tilde{\mathbf{x}}_c$ from class c , we sample independently from each source restricted to the samples of class c yielding $\tilde{\mathbf{s}}_c^{task}$ and mix the data: $\tilde{\mathbf{x}}_c = (\mathbf{W}^{task})^\dagger \tilde{\mathbf{s}}_c^{task}$.

The *Covariance* method generates a new sample of synthetic data in class c by sampling from a Multivariate Gaussian with mean $\boldsymbol{\mu}_c$ and covariance $\boldsymbol{\Sigma}$, where

μ_c is the class mean and Σ is the covariance of centered task data estimated using Ledoit-Wolf method. In brief, it assumes normality of the data per class.

The *ICA + Covariance* method combines the augmentation methods *ICA* and *Covariance*: samples are drawn following the ICA approach, but with some additive non-isotropic Gaussian noise. As in *ICA* we estimate \mathbf{W}^{task} and \mathbf{S}^{task} from \mathbf{X}^{task} via ICA. Then we consider $\mathbf{R}_{task} = \mathbf{X}_{task} - \mathbf{W}_{task}\mathbf{S}_{task}$ and estimate the covariance Σ_R of \mathbf{R}_{task} via LedoitWolf’s method. We then generate a data sample $\tilde{\mathbf{x}}_c$ from class c as with ICA and add Gaussian noise $\tilde{\mathbf{n}} \sim \mathcal{N}(0, \Sigma_R)$. Samples are thus generated as $\tilde{\mathbf{x}}_c + \tilde{\mathbf{n}}$.

The *GANs* (respectively *CGANs*) method use a GANs (respectively CGANs) to generate data. The generator and discriminator have a mirrored architecture with 2 fully connected hidden layer of size (256 and 512). The number of epochs, batch size, momentum and learning rate are set to 20k, 16, 0.9, 0.01 and we use the Leaky RELU activation function.

We evaluate the performance of augmentation methods through the use of classifiers: logistic regression (LogReg), linear discriminant analysis with Ledoit-Wold estimate of covariance (LDA) perceptron with two hidden layers (MLP) and random forrests (RF). The hyper-parameters in each classifier are optimized through an internal 5-Fold cross validation. We set the number of iterations in each classifier so that convergence is reached. The exact specifications are available in appendix Table 4.

4.2 Distinguish fake from real HCP resting state data

This experiment is meant to assess the effectiveness of the data augmentation scheme in producing good samples. Data augmentation methods are trained on 200 subjects taken from HCP rest fMRI dataset which amounts to 960k samples (4800 per individual). Then synthetic data corresponding to 200 synthetic subjects are produced, yielding 960k fake samples and various classifiers are trained to distinguish fake from real data using 5-Fold cross validation. The cross-validated accuracy is shown in Table 1. Interestingly, we observe a dissociation between linear models (LogReg and LDA) that fail to discriminate between generated and actual data, and non-linear models (MLP and RF) that can discriminate samples from alternative augmentation methods. By contrast, all classifiers are at chance when Conditional ICA is used.

4.3 Comparing augmentation methods based on classification accuracy on task HCP dataset

In order to compare the different augmentation methods, we measure their relative benefit in the context of multi-class classification. We use 787 subjects from the HCP task dataset that contains 23 classes and randomly split the dataset into a train set that contains 100 subjects and a test set that contains 687 subjects. In each split we train augmentation methods on the train set to generate fake samples corresponding to 200 subjects. These samples are then appended to the train set, resulting in an augmented train set on which the classifiers are

Models	LDA	LogReg	Random Forest	MLP
ICA	0.493	0.500	0.672	0.697
Covariance	0.473	0.461	0.610	0.626
ICA + Covariance	0.509	0.495	0.685	0.706
GANs [8]	0.501	0.498	0.592	0.607
CGANs [17]	0.498	0.493	0.579	0.604
Conditional ICA	0.503	0.489	0.512	0.523

Table 1: **Distinguish fake from real HCP resting state data** We use HCP resting state data from $n = 200$ subjects (960k samples) and produce an equally large amount of fake data (960k samples) using data augmentation methods. The table shows the 5-fold cross validated accuracy obtained with various classifiers. When Conditional ICA is used, all classifiers are at chance.

trained. Results, displayed in Table 2, show that Conditional ICA always yields a higher accuracy than when no augmentation method is applied. The gains are over 1% on all classifiers tested. By contrast, ICA+Covariance and ICA lead to a decrease in accuracy while the Covariance approach leads to non-significant gains. To further investigate the significance of differences between the proposed approach and other state-of-the-art methods, we perform a t-test for paired samples (see Table 5 in appendix). Most notably, the proposed method performs significantly better than other data augmentation techniques. Given the large size of the HCP task data evaluation set, this significance test would demonstrate that the gains are robust. Note that the random forest classifier is not used in this experiment as it leads to a much lower accuracy than other methods. For completeness, we display the results obtained with random forest in appendix Table 6.

Visualization of fake examples produced by GANs, CGANs and Conditional ICA are available in appendix Figure 5.

4.4 Gains in accuracy brought by conditional ICA on eight datasets.

In this experiment we assess the gains brought by Conditional ICA data augmentation on eight different task fMRI dataset. The number of subjects, classes and the size of the training and test sets differ between dataset and are reported in appendix Table 4. The rest of the experimental pipeline is exactly the same as with the HCP task dataset. We report in Fig. 3 the cross-validated accuracy of classifiers with and without augmentation. We notice that the effect of data augmentation is consistent across datasets, classifiers and splits, with 1% to 5% net gains. An additional experiment studying the sensitivity of Conditional ICA to the number of components used is described in appendix Figure 4.

Models	LDA			LogReg			MLP		
	Acc	Pre	Rec	Acc	Pre	Rec	Acc	Pre	Rec
Original	0.893	0.889	0.891	0.874	0.869	0.873	0.779	0.782	0.778
ICA	0.814	0.809	0.813	0.840	0.836	0.839	0.803	0.805	0.802
Covariance	0.895	0.894	0.895	0.876	0.877	0.875	0.819	0.823	0.820
ICA + Covariance	0.816	0.811	0.812	0.840	0.839	0.840	0.815	0.819	0.814
GANs [8]	0.877	0.871	0.870	0.863	0.865	0.864	0.771	0.779	0.775
CGANs [17]	0.874	0.875	0.872	0.874	0.872	0.875	0.726	0.731	0.725
Conditional ICA	0.901	0.903	0.905	0.890	0.888	0.890	0.832	0.835	0.831

Table 2: **Comparing augmentation methods based on classification accuracy on task HCP dataset** We compare augmentation methods based on the classification accuracy (**Acc**), precision (**Pre**) and recall (**Rec**) obtained by 2 linear classifiers (LDA and LogReg) and one non-linear classifier trained on augmented datasets on HCP Task fMRI data. We report the mean accuracy across 5 splits.

5 Discussion & Future work

In this work we introduced Conditional ICA a fast generative model for rest and task fMRI data. It produces samples that cannot be distinguished from true actual rest by linear as well as non-linear classifiers, showing that it captures higher-order statistics than naive ICA-based generators. When Conditional ICA is used as a data augmentation method, it yields consistent improvement in classification accuracy: on 8 tasks fMRI datasets, we observe increase in accuracy

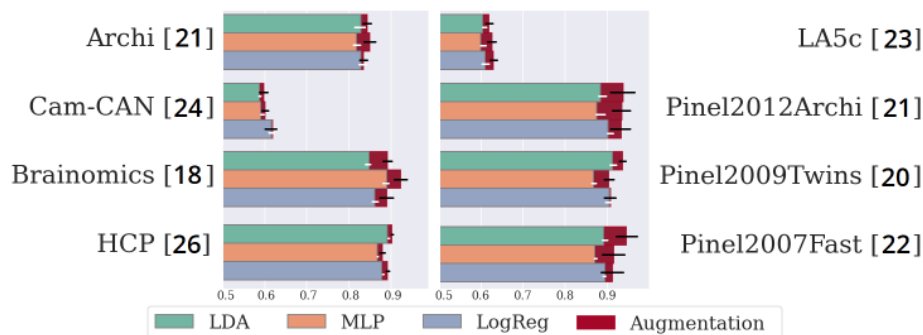


Fig. 3: **Accuracy of models for eight multi-contrast datasets.** Cross validated accuracy of two linear (LDA and LogReg) and one non-linear classifier (MLP) with or without using data augmentation. The improvement yielded by data augmentation is displayed in red. Black error bars indicate standard deviation across splits while white error bars indicate standard deviation across splits with no augmentation.

between 1% and 5% depending on the dataset and the classifier used. Importantly, this performance was obtained without any fine-tuning of the method, showing its reliability. One can also notice that our experiments cover datasets with different cardinalities, from tens to thousand, and different baseline prediction accuracy. It is noteworthy that Conditional ICA is essentially a linear generative model with pointwise non-linearity, which makes it cheap, easy to instantiate on new data, and to introspect.

The speed and simplicity of Conditional ICA as well as the systematic performance improvement it yields, also make it a promising candidate for data augmentation in a wide range of contexts. Future work may focus on its applicability to other decoding tasks such as the diagnosis of Autism Spectrum Disorder (ASD) [4–6] or Attention-Deficit/Hyperactivity Disorder detection (ADHD) [16]. Other extensions of the present work concern the adaptation to individual feature prediction (e.g. age) where fMRI has shown some potential.

Acknowledgements

This project has received funding from the European Union’s Horizon 2020 Framework Program for Research and Innovation under Grant Agreement No 945539 (Human Brain Project SGA3) and the KARAIB AI chair (ANR- 20-CHIA-0025-01).

References

1. Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., Varoquaux, G.: Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics* **8** (2014)
2. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein GAN. arXiv:1701.07875 [cs, stat] (2017)
3. Dadi, K., Varoquaux, G., Machlouzarides-Shalit, A., Gorgolewski, K.J., Wassermann, D., Thirion, B., Mensch, A.: Fine-grain atlases of functional modes for fmri analysis. *NeuroImage* **221**, 117126 (2020)
4. Dvornek, N.C., Ventola, P., Pelphrey, K.A., Duncan, J.S.: Identifying autism from resting-state fmri using long short-term memory networks. In: *International Workshop on Machine Learning in Medical Imaging*. pp. 362–370. Springer (2017)
5. Eslami, T., Mirjalili, V., Fong, A., Laird, A.R., Saeed, F.: Asd-diagnet: a hybrid learning approach for detection of autism spectrum disorder using fmri data. *Frontiers in neuroinformatics* **13**, 70 (2019)
6. Eslami, T., Saeed, F.: Auto-asd-network: a technique based on deep learning and support vector machines for diagnosing autism spectrum disorder using fmri data. In: *Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics*. pp. 646–651 (2019)
7. Friston, K.J., Holmes, A.P., Poline, J., Grasby, P., Williams, S., Frackowiak, R.S., Turner, R.: Analysis of fmri time-series revisited. *Neuroimage* **2**(1), 45–53 (1995)
8. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in neural information processing systems*. pp. 2672–2680 (2014)

9. Gorgolewski, K.J., Varoquaux, G., Rivera, G., Schwarz, Y., Ghosh, S.S., Maumet, C., Sochat, V.V., Nichols, T.E., Poldrack, R.A., Poline, J.B., et al.: Neurovault.org: a web-based repository for collecting and sharing unthresholded statistical maps of the human brain. *Frontiers in neuroinformatics* **9**, 8 (2015)
10. Harris, C.R., Millman, K.J., van der Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., et al.: Array programming with numpy. *Nature* **585**(7825), 357–362 (2020)
11. Hyvärinen, A.: Fast and robust fixed-point algorithms for independent component analysis. *Ieee Trans. Neural Netw* **10**(3), 626–634 (1999)
12. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
13. Kobyzev, I., Prince, S., Brubaker, M.: Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence* p. 1–1 (2020). <https://doi.org/10.1109/tpami.2020.2992934>, <http://dx.doi.org/10.1109/TPAMI.2020.2992934>
14. Ledoit, O., Wolf, M.: A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis* **88**(2), 365–411 (2004)
15. Loula, J., Varoquaux, G., Thirion, B.: Decoding fMRI activity in the time domain improves classification performance. *NeuroImage* **180**, 203–210 (2018)
16. Mao, Z., Su, Y., Xu, G., Wang, X., Huang, Y., Yue, W., Sun, L., Xiong, N.: Spatio-temporal deep learning method for adhd fmri classification. *Information Sciences* **499**, 1–11 (2019)
17. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)
18. Orfanos, D.P., Michel, V., Schwartz, Y., Pinel, P., Moreno, A., Le Bihan, D., Frouin, V.: The brainomics/localizer database. *Neuroimage* **144**, 309–314 (2017)
19. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al.: Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* **12**, 2825–2830 (2011)
20. Pinel, P., Dehaene, S.: Genetic and environmental contributions to brain activation during calculation. *Neuroimage* **81**, 306–316 (2013)
21. Pinel, P., d’Arc, B.F., Dehaene, S., Bourgeron, T., Thirion, B., Le Bihan, D., Poupon, C.: The functional database of the archi project: Potential and perspectives. *NeuroImage* **197**, 527–543 (2019)
22. Pinel, P., Thirion, B., Meriaux, S., Jobert, A., Serres, J., Le Bihan, D., Poline, J.B., Dehaene, S.: Fast reproducible identification and large-scale databasing of individual functional cognitive networks. *BMC neuroscience* **8**(1), 91 (2007)
23. Poldrack, R.A., Congdon, E., Triplett, W., Gorgolewski, K., Karlsgodt, K., Mumford, J., Sabb, F., Freimer, N., London, E., Cannon, T., et al.: A phenome-wide examination of neural and cognitive function. *Scientific data* **3**(1), 1–12 (2016)
24. Shafto, M.A., Tyler, L.K., Dixon, M., Taylor, J.R., Rowe, J.B., Cusack, R., Calder, A.J., Marslen-Wilson, W.D., Duncan, J., Dalgleish, T., et al.: The cambridge centre for ageing and neuroscience (cam-can) study protocol: a cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing. *BMC neurology* **14**(1), 204 (2014)
25. Shirer, W.R., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M.D.: Decoding Subject-Driven Cognitive States with Whole-Brain Connectivity Patterns. *Cerebral Cortex (New York, NY)* **22**(1), 158–165 (2012)
26. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.M.H., et al.: The wu-minn human connectome project: an overview. *Neuroimage* **80**, 62–79 (2013)

27. Varoquaux, G., Thirion, B.: How machine learning is shaping cognitive neuroimaging. *GigaScience* **3**, 28 (2014)
28. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020). <https://doi.org/10.1038/s41592-019-0686-2>
29. Zhuang, P., Schwing, A.G., Koyejo, O.: Fmri data augmentation via synthesis. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). pp. 1783–1787. IEEE (2019)

Dataset	Subjects, Classes	Train/Test	Neurovault collection
hcp [26]	787, 23	100/687	collection 4337
cam-can [24]	605, 5	100/505	collection 4342
brainomics [18]	94, 19	50/44	collection 4341
archi [21]	78, 30	40/38	collection 4339
la5c [23]	191, 24	100/91	collection 4343
pinel2012archi [21]	76, 10	40/36	collection 1952
pinel2009twins [20]	65, 12	35/30	collection 1952
pinel2007fast [22]	133, 10	70/63	collection 1952

Table 3: **Datasets used in the experiments.** The table provides references to the datasets that were used for our experiments, with the number of subjects, the number of classes, the number of subjects in train and test set in each cross validation split and the collection number in Neurovault.

Methods	Optimizer	Hyper-parameters
LogReg	L-BFGS (20 000 iterations)	inverse L_2 regularization strength in $\{0.0001, 0.001, 0.01, 0.1, 1\}$
LDA	Least-squares solver	Estimation of covariance using Ledoit-Wolf’s method
RF	-	Default parameters in sklearn
MLP	Adam (20 000 iterations, momentum: 0.9, batch size: 32, learning rate: 0.0001)	<i>ReLU</i> activation function, fully con- nected architecture with two hidden lay- ers both of size 1024, L2 penalty coeffi- cient: 10^{-5}

Table 4: **Optimizers and hyper-parameters of classifiers** For each classifier, we give the optimization method used as well as the value of hyper-parameters.

Models	LDA	LogR	MLP
Original	$P \leq 0.01$	$P \leq 0.05$	$P \leq 0.0001$
ICA	$P \leq 0.0001$	$P \leq 0.001$	$P \leq 0.001$
Covariance	$P \leq 0.05$	$P \leq 0.01$	$P \leq 0.05$
ICA + Covariance	$P \leq 0.0001$	$P \leq 0.0001$	$P \leq 0.01$
GANs	$P \leq 0.001$	$P \leq 0.001$	$P \leq 0.0001$
CGANs	$P \leq 0.001$	$P \leq 0.01$	$P \leq 0.0001$

Table 5: **Statistical significance of the differences between Conditional ICA and other augmentation methods** P-values obtained after a t-test for paired samples is performed on the data used to produce Table 2.

Methods	Original	ICA	COV.	ICA + COV.	GANs	CGANs	Cond. ICA
Random Forest	0.782	0.778	0.780	0.780	0.780	0.779	0.783

Table 6: **Accuracy of Random Forrest** Mean accuracy obtained from data used for producing the Table 2 using a Random Forest classifier according to the chosen augmentation method.

Methods	Running time (secs)
GANs	12948.2 (\approx 3,60 hr)
CGANs	11015.1 (\approx 3,05 hr)
Conditional ICA	62 s

Table 7: **Running time.** We display the running time of three methods used to generate synthetic data. Conditional ICA is several orders of magnitude faster than GANs or CGANs. In practice the computational overhead induced by Conditional ICA is negligible.




Fig. 4: **Accuracy of augmented discriminative models when varying k .** We use 100 train subjects from the HCP task dataset to train Conditional ICA with k components and generate 200 fake subjects. The classifiers are trained on the train and fake subjects and tested on the left-out 687 subjects. We repeat the procedure for various values of k using 5 random splits per value and report the mean accuracy across splits as a function of k . The dotted line represents the number of components that has been used in our experiments ($k = 900$).

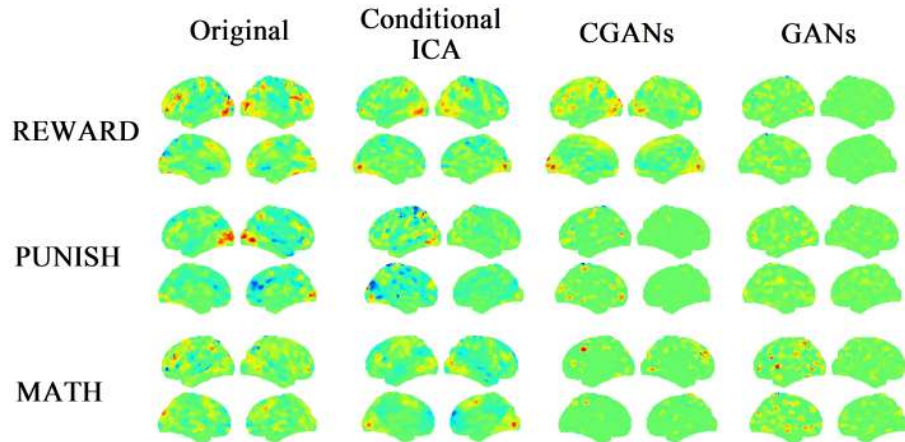


Fig. 5: **Data generation visualization.** Visualization of real (Original) and synthetic brain maps from three generation methods: Conditional ICA (Ours), CGANs and GANs. Three cognitive tasks are shown (reward, punish and math).