

Article

Evaluation of Augmentation Methods in Classifying Autism Spectrum Disorders from fMRI Data with 3D Convolutional Neural Networks

Johan Jönemo ^{1,2}, David Abramian ^{1,2} and Anders Eklund ^{1,2,3,*}

¹ Division of Medical Informatics, Department of Biomedical Engineering, Linköping University, 581 83 Linköping, Sweden

² Center for Medical Image Science and Visualization (CMIV), Linköping University, 581 83 Linköping, Sweden

³ Division of Statistics and Machine Learning, Department of Computer and Information Science, Linköping University, 581 83 Linköping, Sweden

* Correspondence: anders eklund@liu.se

Abstract: Classifying subjects as healthy or diseased using neuroimaging data has gained a lot of attention during the last 10 years, and recently, different deep learning approaches have been used. Despite this fact, there has not been any investigation regarding how 3D augmentation can help to create larger datasets, required to train deep networks with millions of parameters. In this study, deep learning was applied to derivatives from resting state functional MRI data, to investigate how different 3D augmentation techniques affect the test accuracy. Specifically, resting state derivatives from 1112 subjects in ABIDE (Autism Brain Imaging Data Exchange) preprocessed were used to train a 3D convolutional neural network (CNN) to classify each subject according to presence or absence of autism spectrum disorder. The results show that augmentation only provide minor improvements to the test accuracy.

Keywords: functional MRI; resting state; deep learning; augmentation; autism



Citation: Jönemo, J.; Abramian, D.; Eklund, A. Evaluation of Augmentation Methods in Classifying Autism Spectrum Disorders from fMRI Data with 3D Convolutional Neural Networks. *Diagnostics* **2023**, *13*, 2773. <https://doi.org/10.3390/diagnostics13172773>

Academic Editor: Daniele Giansanti

Received: 9 August 2023

Revised: 24 August 2023

Accepted: 25 August 2023

Published: 27 August 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Ever since the emergence of magnetic resonance imaging (MRI) in the 1980s, the absence of ionizing radiation and the flexibility of the acquisition procedure have made this an increasingly important imaging modality in the clinical sciences. The lack of contrast between different tissues in the brain and the interference of the mineralized tissue around it when using X-ray techniques make MRI especially useful in neuroimaging.

While a wide variety of neurological conditions can be diagnosed with MRI, psychiatric anomalies have proven illusive to detect. Presumably, this is because these affect many systems distributed throughout the brain and their manifestations are likely subtle as well as time variant. Furthermore, psychiatric anomalies can vary a lot between subjects. Functional MRI (fMRI) is a technique that seems particularly suited to capture this information, as it generates rich 4D data which can be used for studying brain activity as well as brain connectivity. In this work, it is investigated if deep-learning-based diagnosis of autism from resting state fMRI data can be further improved using 3D augmentation.

1.1. Resting State fMRI

Resting state fMRI has since 1995 been used to study brain connectivity [1,2]. A major advantage compared to task fMRI is that subjects can simply rest during the whole experiment, which normally takes 5–10 min (resulting in some 150–600 brain volumes, or put differently some 50,000 time series), instead of performing different tasks such as finger tapping or mental calculations. This makes it possible to include subjects which for some reason cannot perform certain tasks. A simple measure of the connectivity between two

locations in the brain, called functional connectivity, is the correlation between the two corresponding time series, but several more advanced methods also exist. To limit the size of the 2D correlation matrix, the correlations are normally calculated between the mean time series of some 100–200 brain parcels (instead of some 50,000 voxels). The brain can be divided according to different (resting state) networks, such as the default mode network and the auditory network, and different diseases often affect specific networks.

1.2. Autism

Autism spectrum disorder (ASD) is a disorder characterized by certain features in social communication, and restricted, repetitive, or unusual sensory–motor behaviours [3]. The prevalence of ASD is 1–5% in developed countries [4]. The subject of autism has been studied extensively in recent years, and technology has already contributed to the development of treatments for autism, in terms of rehabilitation and communication.

Due to the lack of reliable biomarkers, the diagnosis is usually based on behaviour, which is very time consuming. Recent work has demonstrated that motor abnormalities can be very informative for detection of ASD [5,6], and that machine learning can be used to shorten the behavioral diagnosis [7]. As ASD results from early altered brain development and neural reorganisation [8,9], it should be possible to derive objective biomarkers from neuroimaging data to aid professionals (paediatricians, psychiatrists, or psychologists) in diagnosing ASD. Here, machine learning can be used to learn informative traits from the high-dimensional fMRI data.

1.3. Machine Learning for Diagnosis of ASD

Several large collaborative efforts have been made to collect and share neuroimaging data of healthy controls as well as diseased [10,11]. ABIDE (Autism Brain Imaging Data Exchange) [12] is one such effort that make available data for 539 subjects diagnosed with ASD as well as 573 typical controls. The ABIDE data originate from 17 sites, and the subjects were aged 7–64 years (median 14.7 years across groups). Using machine learning in an endeavour to classify (resting state) fMRI data according to the presence or absence of ASD has become increasingly popular recently. This classification can be performed in several ways, either using estimated functional connectivity network matrices (2D) or using derivatives (3D volumes), such as weighted and binarized degree centrality, as different approaches to compress the 4D fMRI data. In this work, 3D volumes are used, as it is not obvious how to augment network matrices.

The ASD classification problem seems hard in that accuracies seldom rise to more than 70% when the model classifies unseen data [13–18]. While 1112 subjects is a very large fMRI dataset, it is still small from a deep learning perspective (for example, the popular ImageNet database [19] contains several million images). To further increase the size of the training dataset, and to make convolutional neural networks (CNNs) robust to transformations such as rotation, data augmentation is often used [20,21]. In previous work, it was demonstrated that 3D augmentation for brain tumor segmentation significantly improves the segmentation accuracy [22]. In this work, the purpose is instead to see if 3D augmentation can help train a better ASD classifier, as well as what kind of augmentation techniques work the best.

1.4. Related Work

Several other researchers have used the same ABIDE dataset to train deep learning models for classification [16–18,23,24], but do not mention anything about augmentation. In a recent review on deep learning for autism by Khodatars et al. [25], only advanced augmentation techniques, such as generative adversarial methods (GANs), are briefly mentioned, but training a GAN requires a very large dataset to start from and there is very little work published on 3D GANs. Some researchers have employed resampling techniques wherein shorter time series have been cropped out of longer ones [13,14], typically for the double purpose of getting an augmented data set while also eliminating the extra

complication of variable length sequences. Ji et al. [26] instead applied augmentation to the estimated network matrices. In our study, by contrast, different preprocessing pipelines are used to extract all relevant information from the time dimension, and manipulate data only in the spatial domain.

2. Materials and Methods

2.1. Data

Preprocessing of 4D resting state fMRI data is a complex process involving many different steps, and there is no consensus regarding what the optimal pipeline or toolbox is [27]. Head motion is a major problem in resting state fMRI, as it can, for example, result in erroneous group differences if two cohorts differ in the mean amount of head motion [28,29]. All processing pipelines therefore perform head motion correction, and use additional steps to further suppress motion related signal. ABIDE preprocessed [30] (<http://preprocessed-connectomes-project.org/abide/>, accessed on 10 February 2023) shares preprocessed ABIDE [12] data from structural MRI and resting state fMRI in various forms. As all the preprocessing has been completed, the focus in this work is on the machine-learning-based diagnosis, and other researchers can use the same preprocessed data to reproduce the presented findings. Resting state derivatives (3D volumes where the time dimension has been collapsed into different forms of statistics) resulting from two pipelines were downloaded from ABIDE preprocessed, for 1112 subjects.

One pipeline was the connectome computation system (CCS) [31], which performs slice timing correction, motion realignment, and global intensity normalisation. The data were cleaned from confounders by performing regression with the estimated head movement parameters, the time-dependent global mean intensity, as well as regressors for linear and quadratic drift. Each time series was also band pass filtered (0.01–0.1 Hz). This preprocessing corresponds to the strategy called `global_filt`. Each subject was, furthermore, registered to the MNI152 brain template using boundary based rigid body registration [32] for functional to anatomical registration, and FLIRT and FNIRT for anatomical to template registration [33].

Another such pipeline was “data processing assistant for resting-state fMRI” (DPARSF) [34]. It also performs slice timing correction and motion realignment, but does not perform any intensity normalisation. The same confounders are corrected for and the same band pass filtering is performed, whereupon functional to anatomical registration was performed with ordinary rigid body methods and anatomical to MNI152 brain template registration completed using DARTEL [35].

After preliminary testing of the 10 available derivatives available in ABIDE preprocessed (amplitude of low frequency fluctuations (ALFF), weighted and binarized degree centrality, dual regression, weighted and binarized eigenvector centrality, fractional ALFF, local functional connectivity density (LFC), regional homogeneity (REHO), voxel-mirrored homotopic connectivity (VMHC)), the REHO derivative was chosen for comparing different augmentation strategies. Regional homogeneity is a measure of correlation between a voxel’s time series and those of its neighbours [36], based on the non-parametric rank correlation statistic known as Kendall’s Coefficient of Concordance (KCC) [37]. Each derivative volume from the resting state fMRI data has a size of $61 \times 73 \times 61$ voxels (each $3 \times 3 \times 3 \text{ mm}^3$), which is fed into the 3D CNN described below. See Figure 1 for a preprocessed fMRI volume and the REHO derivative from the CCS pipeline, downloaded from ABIDE (https://s3.amazonaws.com/fcp-indi/data/Projects/ABIDE_Initiative/Outputs/ccs/filt_global/func_preproc/OHSU_0050147_func_preproc.nii.gz, accessed on 1 August 2023; https://s3.amazonaws.com/fcp-indi/data/Projects/ABIDE_Initiative/Outputs/ccs/filt_global/reho/OHSU_0050147_reho.nii.gz, accessed on 1 August 2023). The 539 subjects with ASD and the 573 controls were split 70/15/15 into training, validation, and test sets.

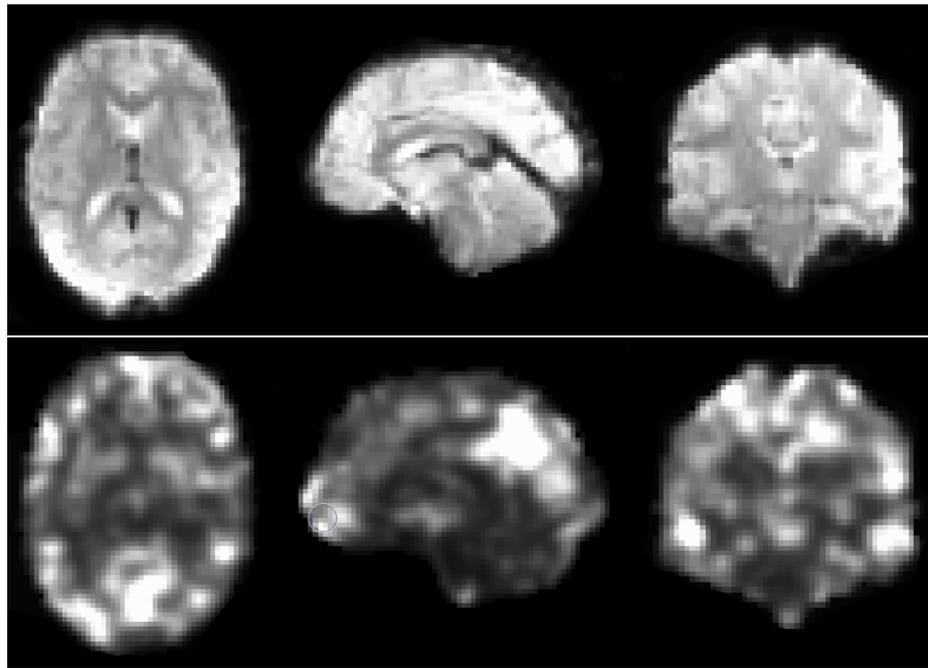


Figure 1. (Top): an fMRI volume obtained after preprocessing with the CCS pipeline. (Bottom): the REHO derivative obtained from the preprocessed 4D fMRI dataset, used by the 3D CNN to classify each subject as control or ASD. Different types of 3D augmentation were applied to each REHO volume, in an attempt to improve the test accuracy. Several other derivatives are available in ABIDE preprocessed, but were not used in this study due to time-consuming training.

2.2. Deep Learning

CNNs are often used for deep-learning-based classification and segmentation of image data, as learning a number of small filters is much more efficient compared to training a dense network (which models the relationship between all pixels in an image, instead of only looking at local correlations). While 2D CNNs are much more common, they are easily extended to 3D as convolution can be performed in any number of dimensions. Unfortunately, existing deep learning frameworks do not support 4D convolutions, which would be required to directly classify 4D fMRI data. The 3D CNN used in this work was implemented using Keras and consists of three convolutional layers (with ReLU activation), max-pooling layers, a dense layer with 16 nodes, and a final one-node layer with sigmoid activation. The first and second convolutional layers contain 8 filters each (size $3 \times 3 \times 3$), and the last convolutional layer uses 16 filters. The total number of trainable parameters in the 3D CNN is approximately 450 k. The CNN was trained with the Adam optimizer with a learning rate of 10^{-5} and a batch size of 16. To prevent overfitting, early stopping was used with a patience of 50 epochs. The training was run until validation accuracy did not improve, and the model was then restored to the state when the last improvement was seen. As an alternative, the models were also trained for 150 epochs with no conditional stopping. To obtain more robust estimates of the test accuracy, 10-fold cross validation was used and the mean test accuracy was calculated.

2.3. Augmentation

There are many types of augmentation that can be useful in 3D. Rotation, flipping, and scaling (zooming in or out) are common for training 2D CNNs, and can also easily be applied in 3D. Elastic (non-linear) deformations are common when training segmentation networks, but perhaps not as common for classification. Brightness augmentation can for example help if the data have been collected at several different MR scanners, as they normally generate data with different brightness [22].

While 2D augmentation functions are included in many deep learning frameworks such as Keras and Pytorch, the support for 3D augmentation is normally lacking. As mentioned by Chlap et al. [21], many researchers use 2D augmentation even if the data are 3D. The 3D augmentation used here is adapted from that of Cirillo et al. [22] and is written in Python/NumPy [38], without facilities for running on a GPU. The 3D augmentation techniques tested in this study are:

- *Flipping*: flipping of the x-axis or not.
- *Rotation*: rotation applied to each axis with angles randomly chosen from a uniform distribution with range between -7.5 and 7.5 degrees, -15 and 15 degrees, -30 and 30 degrees, or -45 and 45 degrees.
- *Scale*: scaling applied to each axis by a factor randomly chosen from a uniform distribution with range $\pm 10\%$ or $\pm 20\%$.
- *Brightness*: power-law γ intensity transformation with its parameters gain (g) and γ chosen randomly between 0.8 and 1.2 from a uniform distribution. The intensity (I) is randomly changed according to the formula: $I_{new} = g \cdot I^\gamma$.
- *Elastic deformation*: elastic (non-linear) deformation with square deformation grid with displacements sampled from from a normal distribution with standard deviation $\sigma = 2, 4, 6, \text{ or } 8$ voxels [39], where the smoothing is done by a spline filter with order 3 in each dimension.

To investigate the effect of combining different types of augmentation, the CNNs were also trained with the two best-performing augmentation approaches according to the CCS pipeline.

The average training time for a single fold were between five minutes and 2.5 h—depending on the type of on-the-fly augmentation employed, the combination of elastic deformation, and an affine transformation being the slowest—using one Nvidia Tesla V100 graphics card for the early stopping models. For the training with a fixed number of epochs, the average single fold training time was at least 10 min but otherwise in the previously mentioned span. In the longer training runs, it is unlikely that the computation speed was bounded by the speed of the graphics card, as the on-the-fly augmentations were performed on the CPU and could be further optimized. In total, some 600 3D CNNs were trained in order to compare all settings.

3. Results

The results from all the different augmentation techniques, as well as baseline results obtained without augmentation, are presented in Figures 2 and 3 (CCS pipeline) and Figures 4 and 5 (DPARSF pipeline). As the dataset is balanced (similar number of ASD and control subjects), only classification accuracy is reported (instead of more advanced metrics, such as area under the curve and Matthew's correlation coefficient). In general, the 3D augmentation does not have a large effect on the test accuracy. For early stopping with the CCS pipeline, random scaling seems to be the best single augmentation approach, but the mean improvement over 10 cross-validation folds is only about 0.5 percentage units. Small elastic deformations also have a small positive effect, while large deformations give worse results.

With the DPARSF pipeline brightness changes appear to be the best augmentation with an increase of 1.9 percentage units, but with high variance over folds, the improvement is negligible. For a fixed number of training epochs, elastic deformations and rotations or combinations thereof seem to work best, with the best improvement of accuracy being 2.2 percentage units in the CCS pipeline and 2.9 percentage units in the DPARSF pipeline. No statistical test was performed to test if this improvement is significant.

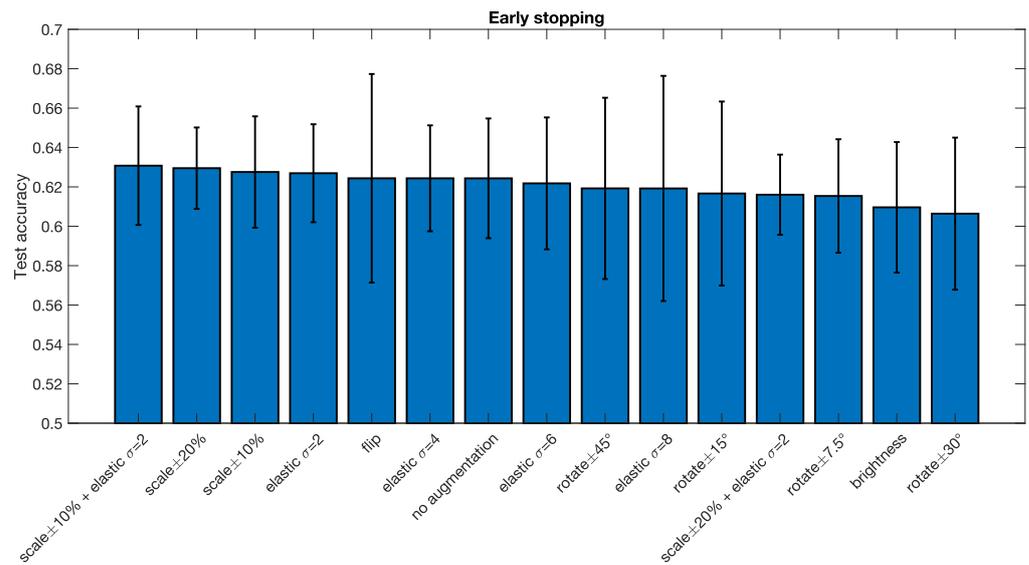


Figure 2. Test accuracy for classifying subjects as healthy or diseased for the ABIDE dataset processed with the CCS pipeline, for different data augmentation approaches. The error bar represents the standard deviation over the 10 cross-validation folds. Note that half of the augmentation approaches result in a test accuracy that is lower compared to the baseline model trained without augmentation, but overall, the differences are small. These results were obtained when using early stopping. Compared to no augmentation, the best augmentation approach increases the test accuracy by 0.6 percentage units.

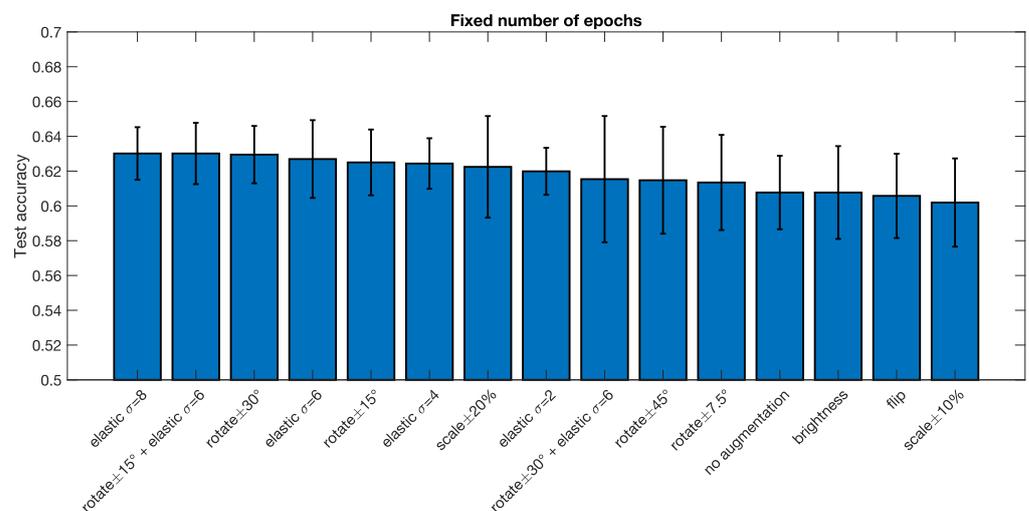


Figure 3. Test accuracy for classifying subjects as healthy or diseased for the ABIDE dataset processed with the CCS pipeline, for different data augmentation approaches. The error bar represents the standard deviation over the 10 cross-validation folds. These results were obtained when using a fixed number of epochs for each training. Compared to no augmentation, the best augmentation approach increases the test accuracy by 2.2 percentage units.

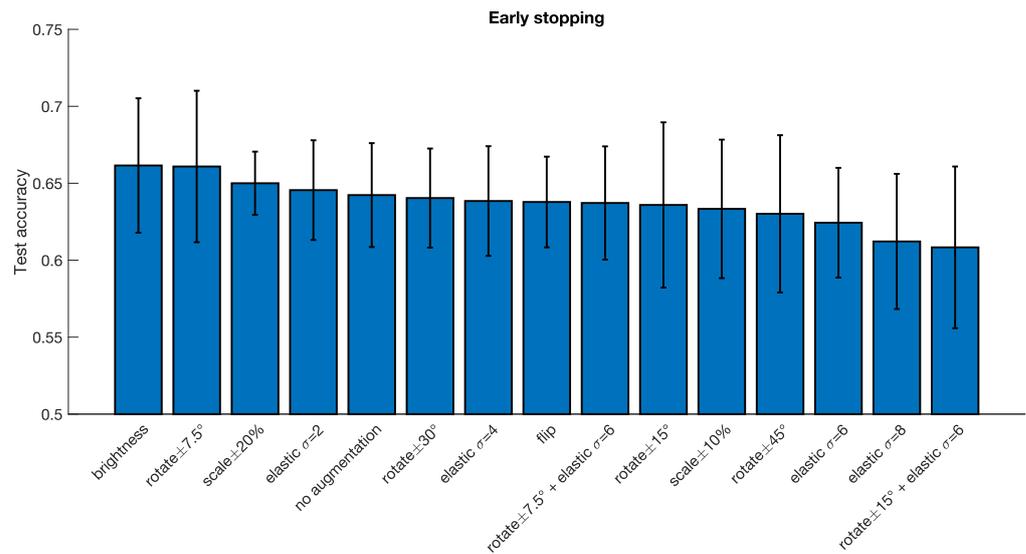


Figure 4. Test accuracy for classifying subjects as healthy or diseased for the ABIDE dataset processed with the DPARSF pipeline, for different data augmentation approaches. The error bar represents the standard deviation over the 10 cross-validation folds. Note that half of the augmentation approaches result in a test accuracy that is lower compared to the baseline model trained without augmentation, but overall, the differences are small. These results were obtained when using early stopping. Compared to no augmentation, the best augmentation approach increases the test accuracy by 1.9 percentage units.

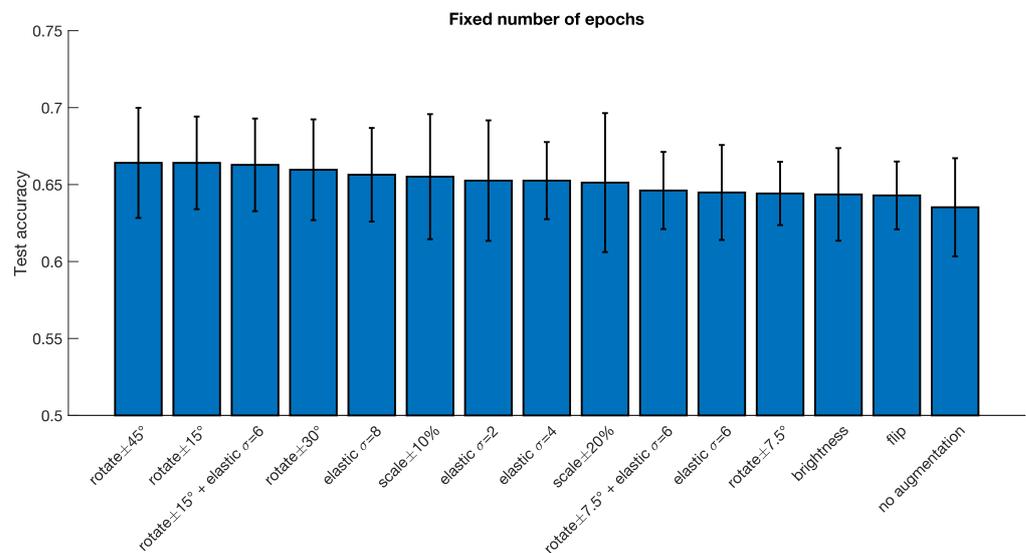


Figure 5. Test accuracy for classifying subjects as healthy or diseased for the ABIDE dataset processed with the DPARSF pipeline, for different data augmentation approaches. The error bar represents the standard deviation over the 10 cross-validation folds. These results were obtained when using a fixed number of epochs for each training. Compared to no augmentation, the best augmentation approach increases the test accuracy by 2.9 percentage units.

4. Discussion

Compared to previous work on 3D augmentation for brain tumor segmentation [22], where several 3D augmentation techniques were shown to significantly improve the segmentation accuracy on the test set, only minor improvements of the test accuracy were found in this study (even though the training accuracy is well above 90%, indicating overfitting). Volume classification is in general a problem which requires more training data compared to volume segmentation, as each volume only represents a single training example, which may partly explain the results.

In this study, brightness augmentation only helps for the DPARSF pipeline with early stopping, while it provided a major improvement for brain tumor segmentation for MR images collected at some 20 different sites [22]. A possible explanation is that the data in this study are not raw MR images, since many preprocessing steps have been used to normalize the intensities to a certain range, and to calculate different derivatives. On the contrary, as the ranges of values in the derivative volumes are not, in general, arbitrary in the same way, brightness augmentation can impair the performance. In DPARSF, no intensity normalization is performed, which may explain why the brightness augmentation results are different compared to the CCS pipeline.

Since all the subjects have been registered to MNI space, it was hypothesized that the results may be different if random transformations are applied to the test volumes, but test time augmentation did not change the findings (results not shown). The presented results are for a single preprocessing strategy (global signal regression and bandpass filtering), and a single derivative, and the preprocessing choice can at least in theory affect how much the augmentation helps.

The focus here has been on classifying ASD and controls, with a binary classifier. ASD criteria are based on DSM-5 criteria, and there are currently three levels of severity. It is possible that using 3D augmentation when training a classifier to distinguish the three severity levels could lead to different results.

The conclusion is that 3D augmentation only provides minor improvements in accuracy (0.6–2.9 percentage units) when training 3D CNNs for classification of ASD versus controls, but the results may be different for an easier task where the baseline test accuracy is for example 80%. The results may also differ for other derivatives in ABIDE preprocessed, and when using several derivatives at the same time using a multi-channel 3D CNN. However, to perform the trainings for many combinations of preprocessing, and for different derivatives, would be very time consuming.

Author Contributions: J.J. and D.A. implemented the 3D CNN and performed the experiments. J.J. and A.E. wrote the paper. A.E. supervised the work. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Swedish research council grant 2017-04889, and by the ITEA/VINNOVA funded project Automation, Surgery Support and Intuitive 3D visualization to optimize workflow in IGT SysTems (ASSIST) (grant 2021-01954), and by the Åke Wiberg foundation (grant M22-0088).

Institutional Review Board Statement: This research study was conducted retrospectively using anonymized human subject data made available in open access by ABIDE preprocessed. Ethical approval was not required as confirmed by the ethics committee of Linköping.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study; see ABIDE for details.

Data Availability Statement: The datasets analyzed for this study can be found at <http://preprocessed-connectomes-project.org/abide/>, accessed on 10 February 2023.

Acknowledgments: Primary support for the work by Adriana Di Martino was provided by the (NIMH K23MH087770) and the Leon Levy Foundation. Primary support for the work by Michael P. Milham and the INDI team was provided by gifts from Joseph P. Healy and the Stavros Niarchos Foundation to the Child Mind Institute, as well as by an NIMH award to MPM (NIMH R03MH096321).

Conflicts of Interest: A.E. has previously received hardware from Nvidia. Otherwise, the authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

1. Biswal, B.; Zerrin Yetkin, F.; Haughton, V.M.; Hyde, J.S. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* **1995**, *34*, 537–541. [[CrossRef](#)] [[PubMed](#)]
2. Van Den Heuvel, M.P.; Pol, H.E.H. Exploring the brain network: A review on resting-state fMRI functional connectivity. *Eur. Neuropsychopharmacol.* **2010**, *20*, 519–534. [[CrossRef](#)] [[PubMed](#)]
3. Lord, C.; Elsabbagh, M.; Baird, G.; Veenstra-Vanderweele, J. Autism spectrum disorder. *Lancet* **2018**, *392*, 508–520. [[CrossRef](#)] [[PubMed](#)]
4. Lyall, K.; Croen, L.; Daniels, J.; Fallin, M.D.; Ladd-Acosta, C.; Lee, B.K.; Park, B.Y.; Snyder, N.W.; Schendel, D.; Volk, H.; et al. The changing epidemiology of autism spectrum disorders. *Annu. Rev. Public Health* **2017**, *38*, 81–102. [[CrossRef](#)] [[PubMed](#)]
5. Simeoli, R.; Milano, N.; Rega, A.; Marocco, D. Using technology to identify children with autism through motor abnormalities. *Front. Psychol.* **2021**, *12*, 635696. [[CrossRef](#)] [[PubMed](#)]
6. Milano, N.; Simeoli, R.; Rega, A.; Marocco, D. A deep learning latent variable model to identify children with autism through motor abnormalities. *Front. Psychol.* **2023**, *14*, 1194760. [[CrossRef](#)] [[PubMed](#)]
7. Wall, D.P.; Dally, R.; Luyster, R.; Jung, J.Y.; DeLuca, T.F. Use of artificial intelligence to shorten the behavioral diagnosis of autism. *PLoS ONE* **2012**, *7*, e43855. [[CrossRef](#)] [[PubMed](#)]
8. Bauman, M.L.; Kemper, T.L. Neuroanatomic observations of the brain in autism: A review and future directions. *Int. J. Dev. Neurosci.* **2005**, *23*, 183–187. [[CrossRef](#)]
9. O'Reilly, C.; Lewis, J.D.; Elsabbagh, M. Is functional brain connectivity atypical in autism? A systematic review of EEG and MEG studies. *PLoS ONE* **2017**, *12*, e0175870. [[CrossRef](#)]
10. Mennes, M.; Biswal, B.B.; Castellanos, F.X.; Milham, M.P. Making data sharing work: The FCP/INDI experience. *Neuroimage* **2013**, *82*, 683–691. [[CrossRef](#)]
11. Poldrack, R.A.; Gorgolewski, K.J. Making big data open: Data sharing in neuroimaging. *Nat. Neurosci.* **2014**, *17*, 1510–1517. [[CrossRef](#)] [[PubMed](#)]
12. 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.; et al. The autism brain imaging data exchange: Towards a large-scale evaluation of the intrinsic brain architecture in autism. *Mol. Psychiatry* **2014**, *19*, 659. [[CrossRef](#)] [[PubMed](#)]
13. Dvornek, N.C.; Ventola, P.; Pelphrey, K.A.; Duncan, J.S. Identifying autism from resting-state fMRI using long short-term memory networks. In *Machine Learning in Medical Imaging: 8th International Workshop, MLMI 2017, Held in Conjunction with MICCAI 2017, Quebec City, QC, Canada, 10 September 2017*; Springer: Berlin/Heidelberg, Germany, 2017; pp. 362–370.
14. Huang, Z.A.; Zhu, Z.; Yau, C.H.; Tan, K.C. Identifying autism spectrum disorder from resting-state fMRI using deep belief network. *IEEE Trans. Neural Netw. Learn. Syst.* **2020**, *32*, 2847–2861. [[CrossRef](#)] [[PubMed](#)]
15. Arbabshirani, M.R.; Plis, S.; Sui, J.; Calhoun, V.D. Single subject prediction of brain disorders in neuroimaging: Promises and pitfalls. *Neuroimage* **2017**, *145*, 137–165. [[CrossRef](#)] [[PubMed](#)]
16. Heinsfeld, A.S.; Franco, A.R.; Craddock, R.C.; Buchweitz, A.; Meneguzzi, F. Identification of autism spectrum disorder using deep learning and the ABIDE dataset. *Neuroimage Clin.* **2018**, *17*, 16–23. [[CrossRef](#)] [[PubMed](#)]
17. Yang, X.; Schrader, P.T.; Zhang, N. A deep neural network study of the ABIDE repository on autism spectrum classification. *Int. J. Adv. Comput. Sci. Appl.* **2020**, *11*, 1–6. [[CrossRef](#)]
18. Thomas, R.M.; Gallo, S.; Cerliani, L.; Zhutovsky, P.; El-Gazzar, A.; van Wingen, G. Classifying autism spectrum disorder using the temporal statistics of resting-state functional MRI data with 3D convolutional neural networks. *Front. Psychiatry* **2020**, *11*, 440. [[CrossRef](#)]
19. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition, Miami, FL, USA, 20–25 June 2009; pp. 248–255.
20. Shorten, C.; Khoshgoftaar, T.M. A survey on image data augmentation for deep learning. *J. Big Data* **2019**, *6*, 60. [[CrossRef](#)]
21. Chlap, P.; Min, H.; Vandenberg, N.; Dowling, J.; Holloway, L.; Haworth, A. A review of medical image data augmentation techniques for deep learning applications. *J. Med. Imaging Radiat. Oncol.* **2021**, *65*, 545–563. [[CrossRef](#)]
22. Cirillo, M.D.; Abramian, D.; Eklund, A. What is the best data augmentation approach for brain tumor segmentation using 3D U-Net? In Proceedings of the IEEE International Conference on Image Processing (ICIP), Anchorage, AK, USA, 19–22 September 2021.
23. Guo, X.; Dominick, K.C.; Minai, A.A.; Li, H.; Erickson, C.A.; Lu, L.J. Diagnosing autism spectrum disorder from brain resting-state functional connectivity patterns using a deep neural network with a novel feature selection method. *Front. Neurosci.* **2017**, *11*, 460. [[CrossRef](#)]
24. El-Gazzar, A.; Quaak, M.; Cerliani, L.; Bloem, P.; van Wingen, G.; Mani Thomas, R. A hybrid 3DCNN and 3DC-LSTM based model for 4D spatio-temporal fMRI data: An ABIDE autism classification study. In *OR 2.0 Context-Aware Operating Theaters and Machine Learning in Clinical Neuroimaging: Second International Workshop, OR 2.0 2019, and Second International Workshop, MLCN 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, 13–17 October 2019*; Springer: Cham, Switzerland, 2019; pp. 95–102.
25. Khodatars, M.; Shoeibi, A.; Sadeghi, D.; Ghaasemi, N.; Jafari, M.; Moridian, P.; Khadem, A.; Alizadehsani, R.; Zare, A.; Kong, Y.; et al. Deep learning for neuroimaging-based diagnosis and rehabilitation of autism spectrum disorder: A review. *Comput. Biol. Med.* **2021**, *139*, 104949. [[CrossRef](#)]
26. Ji, J.; Wang, Z.; Zhang, X.; Li, J. Sparse data augmentation based on encoderforest for brain network classification. *Appl. Intell.* **2022**, *52*, 4317–4329. [[CrossRef](#)]

27. Waheed, S.H.; Mirbagheri, S.; Agarwal, S.; Kamali, A.; Yahyavi-Firouz-Abadi, N.; Chaudhry, A.; DiGianvittorio, M.; Gujar, S.K.; Pillai, J.J.; Sair, H.I. Reporting of resting-state functional magnetic resonance imaging preprocessing methodologies. *Brain Connect.* **2016**, *6*, 663–668. [[CrossRef](#)] [[PubMed](#)]
28. Power, J.D.; Barnes, K.A.; Snyder, A.Z.; Schlaggar, B.L.; Petersen, S.E. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage* **2012**, *59*, 2142–2154. [[CrossRef](#)]
29. Eklund, A.; Nichols, T.E.; Afyouni, S.; Craddock, C. How does group differences in motion scrubbing affect false positives in functional connectivity studies? *BioRxiv* **2020**. [[CrossRef](#)]
30. Craddock, C.; Benhajali, Y.; Chu, C.; Chouinard, F.; Evans, A.; Jakab, A.; Khundrakpam, B.S.; Lewis, J.D.; Li, Q.; Milham, M.; et al. The Neuro Bureau Preprocessing Initiative: Open sharing of preprocessed neuroimaging data and derivatives. *Front. Neuroinform.* **2013**, *7*, 5.
31. Xu, T.; Yang, Z.; Jiang, L.; Xing, X.X.; Zuo, X.N. A connectome computation system for discovery science of brain. *Sci. Bull.* **2015**, *60*, 86–95. [[CrossRef](#)]
32. Greve, D.N.; Fischl, B. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage* **2009**, *48*, 63–72. [[CrossRef](#)]
33. Jenkinson, M.; Beckmann, C.F.; Behrens, T.E.; Woolrich, M.W.; Smith, S.M. FSL. *Neuroimage* **2012**, *62*, 782–790. [[CrossRef](#)]
34. Yan, C.; Zang, Y. DPARSF: A MATLAB toolbox for “pipeline” data analysis of resting-state fMRI. *Front. Syst. Neurosci.* **2010**, *4*, 1377. [[CrossRef](#)]
35. Ashburner, J. A fast diffeomorphic image registration algorithm. *Neuroimage* **2007**, *38*, 95–113. [[CrossRef](#)] [[PubMed](#)]
36. Zang, Y.; Jiang, T.; Lu, Y.; He, Y.; Tian, L. Regional homogeneity approach to fMRI data analysis. *NeuroImage* **2004**, *22*, 394–400. [[CrossRef](#)] [[PubMed](#)]
37. Kendall, M.; Gibbons, J.D. *Rank Correlation Methods*; Oxford University Press: New York, NY, USA, 1990.
38. 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* **2020**, *585*, 357–362. [[CrossRef](#)] [[PubMed](#)]
39. Ronneberger, O.; Fischer, P.; Brox, T. U-net: Convolutional networks for biomedical image segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Munich, Germany, 5–9 October 2015; Springer: Berlin/Heidelberg, Germany, 2015; pp. 234–241.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.