



**HAL**  
open science

# Patch vs. Global Image-Based Unsupervised Anomaly Detection in MR Brain Scans of Early Parkinsonian Patients

Verónica Muñoz-Ramírez, Nicolas Pinon, Florence Forbes, Carole Lartizien, Michel Dojat

► **To cite this version:**

Verónica Muñoz-Ramírez, Nicolas Pinon, Florence Forbes, Carole Lartizien, Michel Dojat. Patch vs. Global Image-Based Unsupervised Anomaly Detection in MR Brain Scans of Early Parkinsonian Patients. MLCN 2021 - 4th International Workshop in Machine Learning in Clinical Neuroimaging, Sep 2021, Strasbourg, France. pp.34-43, 10.1007/978-3-030-87586-2\_4 . hal-03397081

**HAL Id: hal-03397081**

**<https://hal.science/hal-03397081v1>**

Submitted on 22 Oct 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.

# Patch vs. global image-based unsupervised anomaly detection in MR brain scans of early Parkinsonian patients <sup>\*</sup>

Verónica Muñoz-Ramírez<sup>1,2</sup>[0000-0002-0975-7064], Nicolas Pinon<sup>3</sup>[0000-0001-8817-9911] Florence Forbes<sup>2</sup>[0000-0003-3639-0226], Carole Lartizen<sup>3</sup>[0000-0001-7594-4231], and Michel Dojat<sup>1</sup>[0000-0003-2747-6845]

- <sup>1</sup> Univ. Grenoble Alpes, Inserm U1216, CHU Grenoble Alpes, Grenoble Institut des Neurosciences, 38000 Grenoble, France  
{first.last}@univ-grenoble-alpes.fr
- <sup>2</sup> Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France  
{first.last}@inria.fr
- <sup>3</sup> Univ. Lyon, CNRS, Inserm, INSA Lyon, UCBL, CREATIS, UMR5220, U1206, F-69621, Villeurbanne, France  
{first.last}@creatis.insa-lyon.fr

**Abstract.** Although neural networks have proven very successful in a number of medical image analysis applications, their use remains difficult when targeting subtle tasks such as the identification of barely visible brain lesions, especially given the lack of annotated datasets. Good candidate approaches are patch-based unsupervised pipelines which have both the advantage to increase the number of input data and to capture local and fine anomaly patterns distributed in the image, while potential inconveniences are the loss of global structural information. We illustrate this trade-off on Parkinson’s disease (PD) anomaly detection comparing the performance of two anomaly detection models based on a spatial auto-encoder (AE) and an adaptation of a patch-fed siamese auto-encoder (SAE). On average, the SAE model performs better, showing that patches may indeed be advantageous.

**Keywords:** Parkinson’s disease · Anomaly detection · Patches · Siamese networks · Auto-encoders.

## 1 Introduction

Medical imaging represents the largest percentage of data produced in health-care and thus a particular interest has emerged in deep learning (DL) methods

---

<sup>\*</sup> Data used in the preparation of this article were obtained from the Parkinson’s Progression Markers Initiative (PPMI) database ([www.ppmi-info.org/data](http://www.ppmi-info.org/data)).  
· VMR is supported by a grant from NeuroCoG IDEX-UGA (ANR-15-IDEX-02). This work is partially supported by the French program “Investissement d’Avenir” run by the Agence Nationale pour la Recherche (ANR-11-INBS-0006).  
· VMR and NP contributed equally to this work.

to create support tools for radiologists to analyze multimodal medical images, segment lesions and detect subtle pathological changes that even an expert eye can miss. The vast majority of these methods are based on supervised models which require to be trained on large series of annotated data, time- and resource-consuming to generate.

Over the years, several publicly available neuroimaging databases have been curated and completed with annotations. Some of the most prominent ones are: MSSEG, for multiple sclerosis lesion segmentation [4]; BRATS, for brain tumor segmentation [14]; ISLES, for ischemic stroke lesion segmentation [13]; and mTOP for mild traumatic injury outcome prediction [12]. Challenges, namely those of MICCAI, are organized regularly to showcase the latest technological advancements and push the community towards better performances. However, there are several neurological diseases seldom studied due to the small size and subtlety of the lesions they present. This is the case of vascular disease, epilepsy, and most neurodegenerative diseases in their early stages. The main challenge for such pathologies indeed, is to identify the variability of the pathological patterns on images where the lesion is barely seen or not visible.

Unsupervised methods are good candidates to tackle both the lack of annotated examples and the subtlety of brain scan anomalies [3,19]. They rely on networks that learn to encode normal brain patterns in such a manner that any atypical occurrence can be identified by the inability of the network to reproduce it. Auto-encoders (AE), variational auto-encoders (VAE) [10] and generative adversarial networks (GAN) [7], have been extensively used as building blocks for unsupervised anomaly detection due to their ability to learn high-dimensional distributions [3].

Parkinson’s Disease (PD) is a neurodegenerative disorder that is only identifiable through routine MR scans at an advanced stage. Nevertheless, the manifestation of non-motor symptoms, years before the apparition of the first motor disturbances, suggests the presence of physio-pathological differences that could allow for earlier PD diagnosis. PD afflicts patients for as many as one to two decades of their lives and current treatments can only attenuate some motor manifestations [21]. Therefore, reducing the gap between diagnosis and the onset of the neurodegenerative process is of paramount importance to identify personalized treatments that would significantly slow its natural progression. Unsupervised anomaly detection models are here employed to explore such challenging MR data analysis.

In a previous work [16], we compared deterministic and variational, spatial and dense autoencoders for the detection of subtle anomalies in the diffusion parametric maps of *de novo* (i.e., newly diagnosed and without dopaminergic treatment) PD patients from the PPMI database [11]. Our results, while preliminary, offered compelling evidence that DL models are useful to identify subtle anomalies in early PD, even when trained with a moderate number of images and only two parametric maps as input.

Our goal in this paper is to compare an improved anomaly detection pipeline based on a deterministic spatial auto-encoder, hereafter simply referred as AE, to

an adaptation of patch-based siamese auto-encoder (SAE) proposed in [1]. This architecture was originally intended to the detection of subtle epileptic lesions, application for which it achieved promising results.

One important difference between the two compared architectures is the dimension of the input and output data. While AE were trained on 2D transverse slices, thus capturing a global pattern in the image, SAE were trained on small patches sampled throughout the data, making them more suitable to capture fine patterns but losing global structural integrity. Through this comparison we aim to analyze the advantages of patch-fed architectures for the identification of subtle and local abnormalities as well as to propose an alternative for anomaly detection in moderate size image datasets.

## 2 Brain anomaly detection pipeline

The anomaly detection task with auto-encoders can be formally posed as follows:

- An auto-encoder is first trained to reconstruct normal samples as accurately as possible. This network is composed of two parts: an encoder (1) that maps the input data into a lower dimensional latent space, assumed to contain important image features, and a decoder (2) that maps the code from the latent space into an output image.
- When fed by an unseen image, this trained network produces a reconstructed image from its sampled latent distribution which is the counterpart ‘normal’ part of the input image.
- *Reconstruction error* maps, computed as the difference between the input and output images, are thus assumed to highlight anomalous regions of the input data.
- *Anomaly scores* at the voxel, region of interest or image levels may then be derived from the post-processing of these *reconstruction error* maps.


In this work, we present a general framework for unsupervised brain anomaly detection based on auto-encoders to produce reconstruction error maps and a novel post-processing step to derive per-region anomaly scores.

### 2.1 Autoencoder architectures

We constructed and evaluated two auto-encoder models: a classic auto-encoder (AE) and a siamese auto-encoder (SAE). Both models are fully-convolutional. Their architectures are displayed in Figure 1 and their differences are detailed below.

**Classic Auto-encoder:** This architecture consists of 5 convolutional layers that go from input to bottleneck and 5 transposed convolutional layers going from bottleneck to output. The output of the encoder network is directly the latent vector  $z$  and the loss function was simply the  $L_1$ -norm reconstruction error between input  $\mathbf{x}$  and output  $\hat{\mathbf{x}}$ :

$$L_{AE}(\mathbf{x}) = \|\mathbf{x} - \hat{\mathbf{x}}\|_1 \quad (1)$$



**Fig. 1.** Classical auto-encoder (AE) on top, Siamese auto-encoder (SAE) at the bottom

**Siamese Auto-encoder:** The siamese autoencoder (SAE) model [1] consists of two identical convolutional autoencoders with shared parameters. The SAE receives a pair of patches  $(\mathbf{x}_1, \mathbf{x}_2)$  at input that are propagated through the network, yielding representations  $\mathbf{z}_t \in \mathcal{Z}, t = (1, 2)$  in the middle layer bottleneck. The second term of the loss function  $L_{SAE}$  (eq. 2) is designed to maximize the cosine similarity between  $\mathbf{z}_1$  and  $\mathbf{z}_2$ . This constraints patches that are "similar" to be aligned in the latent space. Unlike standard siamese architectures where *similar* and *dissimilar* pairs are presented to the network, Alaverdyan et al. [1] proposed to train this architecture on *similar* pairs only and compensate for the lack of *dissimilar* pairs through a regularizing term that prevents driving the loss function to 0 by mapping all patches to a constant value. This term is defined as the mean squared error between the input patches and their reconstructed outputs. The proposed loss function for a single pair hence is:

$$L_{SAE}(\mathbf{x}_1, \mathbf{x}_2) = \sum_{t=1}^2 \|\mathbf{x}_t - \hat{\mathbf{x}}_t\|_2^2 - \alpha \cdot \cos(\mathbf{z}_1, \mathbf{z}_2) \quad (2)$$

where  $\hat{\mathbf{x}}_t$  is the reconstructed output of the patch  $\mathbf{x}_t$  while  $\mathbf{z}_t$  is its representation in the middle layer bottleneck and  $\alpha$  an hyperparameter that controls the trade-off between the two terms.

As depicted on Figure 1, the encoder part is composed of 3 convolutional layers and one maxpooling layer in-between the first and second convolutions, while the (non symmetrical) decoder part is composed of 4 convolutional layers, with an upsampling layer in-between the second and third convolutional layer.

## 2.2 Post-processing of the reconstruction error maps

We leveraged the reconstruction error maps obtained from both architectures to generate an anomaly score, following the methodology introduced in [15]. The voxel-wise reconstruction errors in one image were computed as  $\|x_i - \hat{x}_i\|_1$ . Since the architectures were fed more than one channel (here two MR modalities), we defined the joint reconstruction error of every voxel as the square root of the sum of squares of the difference between input and output for every channel.

Next, we fixed a threshold on these generated *reconstruction error maps* to decide whether or not a given voxel should be considered as abnormal, hereafter called the *abnormality threshold*. Since we expected PD patients to exhibit abnormal voxels in larger quantities than controls, this value corresponded to an extreme quantile (eg. the 98% quantile) of the errors distribution in the control population. The thresholded reconstruction error maps were then employed to identify anatomical brain regions for which the number of abnormal voxels could be used to discriminate between patients and controls.

## 3 Experiments

### 3.1 Data

The dataset used in this work consisted of DTI MR scans of 57 healthy controls and 129 *de novo* PD patients selected from the PPMI database. All images were acquired with the same MR scanner model (3T Siemens Trio Tim) and configured with the same acquisition parameters. Only one healthy control was taken out of the study due to important artifacts in the images.

From these images, two measures, mean diffusivity (MD) and fractional anisotropy (FA), were computed using MRtrix3.0. Values of FA and MD were normalized into the range  $[0, 1]$ . The images were spatially normalized to the standard brain template of the Montreal Neurological Institute (MNI) with a non-linear deformation. The resulting MD and FA parameter maps were of dimension  $121 \times 145 \times 121$  with a voxel size of  $1.5 \times 1.5 \times 1.5\text{mm}^3$ .

The control dataset was divided into 41 training controls and 15 testing controls to avoid data leakage. This division was effectuated in 10 different manners through a bootstrap procedure in order to assess the generalization of our predictions as advised in [17]. We took special care to maintain an age average around 61 years old for all the training and test population as well as a 40-60 proportion of females and males.

Once the models were trained with one of the 10 training sets, they were evaluated with the corresponding healthy control test set and the PD dataset (age: 62 y.  $\pm$  9; sex: 48 F).

### 3.2 Training of the auto-encoders

**AE models** The training dataset of the AE models consisted of 1640 images corresponding to 40 axial slices around the center of the brain for each of the

41 training control subjects. The AE models were trained for 160 epochs, with a learning rate of  $10^{-3}$ .  $3 \times 3$  kernels were convoluted using padding of 1 pixel and a stride of (2, 2). The bottleneck dimensions were  $h=4$ ,  $w=5$  and  $c=256$ . There were no pooling layers. Implementation was done in Python 3.6.8, PyTorch 1.0.1, CUDA 10.0.130 and trained on a NVIDIA GeForce RTX 2080 Ti GPU with batches of 40 images. After each convolutional layer, batch normalization [8] was applied for its regularization properties. The rectified linear unit (ReLU) activation function was employed in each layer except the last, for which a sigmoid was preferred. The loss functions were optimized using Adam [9].

**SAE models** The SAE model was trained with 600 000 patches of size  $15 \times 15 \times 2$  ( $\sim 15$  000 patches per subject). The model was trained for 30 epochs, with a learning rate of  $1 \times 10^{-3}$ . Bottleneck dimensions were  $h=2$ ,  $w=2$  and  $c=16$ . Maxpooling and upsampling layers were used, as detailed before. No batch-normalization layers were used. Activation function for every convolutional layer was the rectified linear unit (ReLU) and the sigmoid function was used in the last layer. The kernel size and the numbers of filters were  $3 \times 3$  and 16 respectively for all convolution blocks but the final one with  $2 \times 2$  and 2 (equal to the number of channels) respectively. The stride was 1 for all convolution blocks. The maxpooling/upsampling factor was 2. Implementation was done in Python 3.8.10, Tensorflow 2.4.1, 11.0.221 and trained on a NVIDIA GeForce GTX 1660 GPU with batches of 225 patches. The loss function was comprised of a reconstruction part (mean squared error) and a similarity measure (cosine similarity) weighted by a coefficient  $\alpha=0.005$ . The loss function was minimized using Adam [9].


### 3.3 Performance evaluation

The percentage of abnormal voxels found in the thresholded reconstruction error maps was employed to classify them as healthy or pathological (PD) based on a threshold. The critical choice of the threshold was investigated using a Receiver Operator Curve (ROC), taking into account the imbalanced nature of our test set (15 healthy and 129 PD). Every point in the ROC corresponds to the sensitivity and specificity values obtained by a given threshold. As proposed in [15], the choice of the optimal threshold, referred to as the *pathological threshold*, was based on the optimal geometric mean,  $g\text{-mean} = \sqrt{\text{Sensitivity} \times \text{Specificity}}$ .

Additionally, to help evaluate the localization of anomalies, two atlases were considered: the Neuromorphometrics atlas [2] and the MNI PD25 atlas [20]. The first was used to segment the brain into 8 macro-regions: subcortical structures, white matter and the 5 gray matter lobes (Frontal, Temporal, Parietal, Occipital, Cingulate/Insular). The latter was specifically designed for PD patients exploration. It contains 8 regions: substantia nigra (SN), red nucleus (RN), subthalamic nucleus (STN), globus pallidus interna and externa (GPi, GPe), thalamus, putamen and caudate nucleus. For all of the before-mentioned regions of interest (ROI), we calculated the *g-mean* for the associated pathological threshold, leading to the classification performance of our models.


## 4 Results

As it can be seen in Figure 2, both auto-encoder architectures achieve good quality reconstructions, however the SAE seems to capture finer details and textures than the AE. This explains the high contrast in AE reconstructions.



**Fig. 2.** Showcase of a slice of the original data and its AE and SAE reconstructions

The visualization of the percentage of abnormal voxels in the ROIs presented in Figure 3 showcases the inter-subject variability amongst members of the same population (healthy and PD). Even so, abnormal voxels are clearly more numerous in the PD patient population.




**Fig. 3.** The percentage of abnormal voxels found by the SAE in the anatomical ROIs presented in Section 3. Top: the test controls of Sample 1; Bottom: 15 randomly selected PD patients.

The *g-mean classification* scores for all models, obtained for each ROI and each sub-population, are presented in Figure 4. We notice that, on average, the SAE model performed better than the AE on the whole brain and most of the macro and subcortical structures studied, with the exception of the temporal lobe, the putamen, the thalamus and the internal and external segment of the globus pallidus. We note that the results varied greatly across the ten populations samples. As an indication, for the whole brain, the SAE obtained a *g-mean* average score of  $66.9 \pm 5.8\%$  and the AE  $65.3 \pm 7.5\%$ , however the best scores



among the 10 samples were of 79.9% and 81.9% for the AE and SAE respectively, both on sample 1. Corresponding values for the white matter are of  $68.2 \pm 4.6\%$  for SAE and  $66.2 \pm 6.7\%$  for AE. The largest standard deviations in the observed anatomical regions belonged to the white matter, the frontal and occipital lobes.



**Fig. 4.** *g-mean* scores for the whole brain and several ROIs for AE and SAE. The vertical dashed lines separate macro- and micro brain structures.

## 5 Discussion and conclusion

Unsupervised auto-encoders (AE) have shown to efficiently tackle challenging detection tasks where brain alteration are barely seen or not visible. The objective of our study was to explore the potential of such AE models for the detection of subtle anomalies in *de novo* PD data and compare patch-based versus image-based models.

Both AE and SAE architectures produced good quality reconstructions and were able to discriminate between healthy individuals and recently diagnosed PD patients with performances (see Figures 3 and 4) that are competitive with those found in the literature. Notably, the Correia & al. [5] SVM mean accuracy score for a selection of WM regions is of 61.3% whereas both the SAE and AE achieved *g-mean* scores above 66% for the WM. Also, the cross-validation procedure of Schuff & al. [18] obtained a ROC AUC of 59% for the rostral segment of the SN which is below our SN average *g-mean* score for the SAE and equal to that of the AE.

Note that at this early stage of the disease (1-2 H-Y scale) the patients have no tremor nor uncontrolled movements compared to healthy subjects. This rules out that movement was the index that allowed PD classification.

Using DTI data we did not search for structural atrophy or lesion load but for degradation of WM properties in the early stages of PD that could appear everywhere in the brain. This explains why the WM obtains the highest g-mean scores. This being said, our models could largely improve by increasing the size of our dataset. Furthermore, the addition of another MR modality such as iron load using T2/T2\* relaxometry could allow us to detect the reduction of dopaminergic neurons in subcortical structures, largely reported in the early stages of PD but not visible in DTI.

Regarding the comparison between the AE and SAE, the choice is not clear. While the classic AE architecture benefits from a more straightforward implementation, the SAE proposes significant advantages for small databases. Indeed, patch-fed networks can be trained with smaller samples of data and the siamese constraint of the architecture ensures efficient learning. What is more, the latent space features of these models contain local information that can be used to classify between healthy and pathological individuals at the voxel-level and produce anomaly maps like those introduced by [1].

In future work, we plan to generate this kind of maps for early PD patients to offer more precise indications about the localization of anomalies and correlate them with the PD hemispheric lateralization. In addition, we aim to complete our dataset by adding other MR modalities such as perfusion and relaxometry, but also by gathering heterogeneous data from multi-vendors and multi-sites exams. For this purpose, we will use an harmonization procedure as a preprocessing step (extension of DeepHarmony [6]). Finally, our 3D implementation of the SAE model is ongoing.

## References

1. Alaverdyan, Z., Jung, J., Bouet, R., Lartizien, C.: Regularized siamese neural network for unsupervised outlier detection on brain multiparametric magnetic resonance imaging: Application to epilepsy lesion screening. *Medical Image Analysis* **60**, 101618 (2020). <https://doi.org/10.1016/j.media.2019.101618>
2. Bakker, R., Tiesinga, P., Kotter, R.: The Scalable Brain Atlas: Instant Web-Based Access to Public Brain Atlases and Related Content. *Neuroinformatics* **13**, 353–366 (2015)
3. Baur, C., Denner, S., Wiestler, B., Navab, N., Albarqouni, S.: Autoencoders for unsupervised anomaly segmentation in brain MR images: A comparative study. *Medical Image Analysis* **69**, 101952 (2021). <https://doi.org/https://doi.org/10.1016/j.media.2020.101952>
4. Commowick, O., Istace, A., Kain, M., Laurent, B., Leray, F., Simon, M., Pop, S.C., Girard, P., Améli, R., Ferré, J.C., Anne Kerbrat, T.T., Frédéric Cervenansky, T.G., Beaumont, J., Doyle, S., Forbes, F., Knight, J., Khademi, A., Mahbod, A., Wang, C., McKinley, R., Wagner, F., Muschelli, J., Sweeney, E., Roura, E., Lladó, X., Santos, M.M., Santos, W.P., Silva-Filho, A.G., Tomas-Fernandez, X., Hélène Urien, I.B., Valverde, S., Cabezas, M., Vera-Olmos, F.J., Malpica, N., Guttmann, C., Sandra Vukusic, G.E., Dojat, M., Styner, M., Warfield, S.K., Cotton, F., Barillot, C.: Objective evaluation of multiple sclerosis lesion segmentation using a data

- management and processing infrastructure. *scientific reports* 8, 13650 (2018), <http://portal.fli-iam.irisa.fr/msseg-challenge>
5. Correia, M.M., Rittman, T., Barnes, C.L., Coyle-Gilchrist, I.T., Ghosh, B., Hughes, L.E., Rowe, J.B.: Towards accurate and unbiased imaging-based differentiation of Parkinson's disease, progressive supranuclear palsy and corticobasal syndrome. *Brain Communications* (2020)
  6. Dewey, B.E., Zhao, C., Reinhold, J.C., Carass, A., Fitzgerald, K.C., Sotirchos, E.S., Saidha, S., Oh, J., Pham, D.L., Calabresi, P.A., van Zijl, P.C., Prince, J.L.: DeepHarmony: A deep learning approach to contrast harmonization across scanner changes. *Magnetic Resonance Imaging* (Jul 2019). <https://doi.org/10.1016/j.mri.2019.05.041>
  7. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative Adversarial Networks. *arXiv:1406.2661 [cs, stat]* (Jun 2014)
  8. Ioffe, S., Szegedy, C.: Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv:1502.03167 [cs]* (Mar 2015)
  9. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]* (Jan 2017)
  10. Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. *arXiv:1312.6114 [cs, stat]* (May 2014)
  11. Marek, K., Chowdhury, S., Siderowf, A., Lasch, S., Coffey, C., Caspell-Garcia, C., Simuni, T., Jennings, D., M. Tanner, C., Q. Trojanowski, J., Shaw, L., Seibyl, J., Schuff, N., Singleton, A., Kiebertz, K., W. Toga, A., Mollenhauer, B., Galasko, D., M. Chahine, L., Initiative, P.P.M.: The parkinson's progression markers initiative (ppmi) - establishing a pd biomarker cohort. *Annals of Clinical and Translational Neurology* p. 1460–1477 (2018)
  12. MICCAI: Mild traumatic brain injury outcome prediction (2016), [www.tbichallenge.wordpress.com](http://www.tbichallenge.wordpress.com)
  13. MICCAI: Ischemic stroke lesion segmentation challenge (2018), [www.isles-challenge.org](http://www.isles-challenge.org)
  14. MICCAI: Brain tumor segmentation challenge (2020), <http://braintumorsegmentation.org/>
  15. Muñoz-Ramírez, V., Kmetzsch, V., Forbes, F., Meoni, S., Moro, E., Dojat, M.: Subtle anomaly detection in mri brain scans: Application to biomarkers extraction in patients with de novo parkinson's disease. *medRxiv* (2021). <https://doi.org/10.1101/2021.06.03.21258269>
  16. Muñoz Ramírez, V., Kmetzsch, V., Forbes, F., Dojat, M.: Deep learning models to study the early stages of parkinson's disease. In: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI). pp. 1534–1537 (2020). <https://doi.org/10.1109/ISBI45749.2020.9098529>
  17. Poldrack, R.A., Huckins, G., Varoquaux, G.: Establishment of Best Practices for Evidence for Prediction: A Review. *JAMA Psychiatry* pp. 534–540 (2019)
  18. Schuff, N., Wu, I.W., Buckley, S., Foster, E.D., Coffey, C.S., Gitelman, D.R., Mendick, S., Seibyl, J., Simuni, T., Zhang, Y., Jankovic, J., Hunter, C., Tanner, C.M., Rees, L., Factor, S., Berg, D., Wurster, I., Gauss, K., Sprenger, F., Seppi, K., Poewe, W., Mollenhauer, B., Knake, S., Mari, Z., Mccoy, A., Ranola, M., Marek, K.: Diffusion imaging of nigral alterations in early Parkinson's disease with dopaminergic deficits. *Movement Disorders* **30**, 1885–1892 (2015)
  19. Shinde, S., Prasad, S., Saboo, Y., Kaushick, R., Saini, J., Pal, P.K., Ingalthalakar, M.: Predictive markers for Parkinson's disease using deep neural nets on neuromelanin sensitive MRI. *NeuroImage: Clinical* **22**, 101748 (2019)

20. Xiao, Y., Fonov, V., Beriault, S., Al Subaie, F., Chakravarty, M.M., Sadikot, A.F., Pike, G.B., Collins, D.L.: Multi-contrast unbiased MRI atlas of a Parkinson's disease population. *Int J Comput Assist Radiol Surg* **10**, 329–341 (2015)
21. Zhao, Y.J., Wee, H.L., Chan, Y.H., Seah, S.H., Au, W.L., Lau, P.N., Pica, E.C., Li, S.C., Luo, N., Tan, L.C.: Progression of Parkinson's disease as evaluated by Hoehn and Yahr stage transition times. *Movement Disorders* **25**(6), 710–716 (apr 2010). <https://doi.org/10.1002/mds.22875>