

# Interpretable Brain Disease Classification and Relevance-Guided Deep Learning

Christian Tinauer<sup>1,2</sup>, Stefan Heber<sup>1</sup>, Lukas Pirpamer<sup>1</sup>, Anna Damulina<sup>1</sup>, Reinhold Schmidt<sup>1</sup>, Rudolf Stollberger<sup>2,3</sup>, Stefan Ropele<sup>1,3</sup>, and Christian Langkammer<sup>1,3</sup>✉

<sup>1</sup>Department of Neurology, Medical University of Graz, Graz, Austria  
<sup>2</sup>Institute of Medical Engineering, Graz University of Technology, Graz, Austria  
<sup>3</sup>BioTechMed-Graz, Austria

Deep neural networks are increasingly used for neurological disease classification by MRI, but the networks' decisions are not easily interpretable by humans. Heat mapping by deep Taylor decomposition revealed that (potentially misleading) image features even outside of the brain tissue are crucial for the classifier's decision. We propose a regularization technique to train convolutional neural network (CNN) classifiers utilizing relevance-guided heat maps calculated online during training. The method was applied using T1-weighted MR images from 128 subjects with Alzheimer's disease (mean age=71.9±8.5 years) and 290 control subjects (mean age=71.3±6.4 years). The developed relevance-guided framework achieves higher classification accuracies than conventional CNNs but more importantly, it relies on less but more relevant and physiological plausible voxels within brain tissue. Additionally, preprocessing effects from skull stripping and registration are mitigated. With the interpretability of the decision mechanisms underlying CNNs, these results challenge the notion that unprocessed T1-weighted brain MR images in standard CNNs yield higher classification accuracy in Alzheimer's disease than solely atrophy.

Deep convolutional neural networks | heat mapping | relevance guidance | Alzheimer's disease

Correspondence: [christian.langkammer@medunigraz.at](mailto:christian.langkammer@medunigraz.at), Department of Neurology, Medical University of Graz, Auenbruggerplatz 22, 8036 Graz, Austria

## Introduction

Alzheimer's disease (AD) is the most common form of dementia with about 50 million patients and a substantial burden for our healthcare systems, caregivers and next of kin (Scheltens et al., 2021). While postmortem diagnosis can be obtained from the histological examination of tissue samples from affected anatomical regions (Braak et al., 2006; Braak and Braak, 1991), in vivo diagnosis is hampered by clinical symptom similarities and its accuracy is rather low (71%–87% sensitivity and 44%–71% specificity) (Oldan et al., 2021).

In addition to clinical and neuropsychological tests, medical imaging is increasingly used to strengthen diagnosis by PET imaging ligands to amyloid- $\beta$  and tau proteins combined with MRI. Recently revised diagnosis criteria for AD are clinical-biological and require both clinical phenotype and biomarker evidence (A $\beta$  or tau) of AD (Dubois et al., 2021). Although the presence of extracellular neuritic A $\beta$  plaques is part of several diagnosis criteria their clinical value is discussed con-

troversially, whereas selective tau ligands do reflect clinical severity and memory impairment and also serve for in-vivo Braak-staging (Biel et al., 2021). Based on imaging tau pathology, recent fascinating data-driven work found that tau-PET can be used to identify four spatiotemporal phenotypes which exhibit different clinical profiles and longitudinal outcomes and thus opens an avenue for personalized treatment (Vogel et al., 2021). However, AD has a long prodromal and asymptomatic inflammatory phase where radioactive PET tracers cannot be used as a means for its prognosis in a healthy population. Because pathological changes are occurring decades before initial clinical manifestations, early biomarkers in a broad population might be obtained best by MRI, where volumetry and especially hippocampal atrophy are presently used as imaging markers (Henneman et al., 2009; Leung et al., 2013; Sluimer et al., 2008).

Deep learning is omnipresent in medical imaging, including image reconstruction (Hammernik et al., 2018), segmentation (Kleesiek et al., 2016), and classification (Esteva et al., 2017; Bäckström et al., 2018). Convolutional neural networks (CNNs) are utilized for neurological disease classification (Noor et al., 2020; Vieira et al., 2017; Zhang et al., 2020) and regression (Dinsdale et al., 2021a) in prevalent neurological disorders such as Alzheimer's disease (Oh et al., 2019; Bäckström et al., 2018; Böhle et al., 2019; Korolev et al., 2017), Parkinson's disease (Karapinar Senturk, 2020) and multiple sclerosis (Eitel et al., 2019).

Despite their improved performance, those models are generally not easily interpretable by humans and deep neural networks (DNNs) are mostly seen as black boxes where data in combination with extensive learning efforts yields decisions (Davatzikos, 2019). One striking example of misguided feature extraction of DNNs is described in (Lapuschkin et al., 2019), where secondary photo watermarks identified horses better than the actual animal print. In the context of brain MRI it has been shown that learned features for age estimation are influenced by the applied registration type (linear vs. nonlinear) (Dinsdale et al., 2021a). However, no systematic investigation of the preprocessing of brain MR images for disease classification with CNNs has been conducted, but the studies (Böhle et al., 2019; Eitel et al., 2019; Oh et al., 2019) aimed at explaining their applied classifier. Preprocessing is a crucial step, with skull stripping (brain extraction) creating artificial edges and interpolation and regridding necessary for registration. CNNs can incorporate these newly introduced

**NOTE: This preprint reports new research that has not been certified by peer review and should not be used to guide clinical practice.**

features during training and base their classification results thereon.

Medical imaging has high legal requirements as e.g. the EU's General Data Protection Regulation (GDPR) explicitly requires the right to explanation for users subjected to decisions of an automated processing system (Goodman and Flaxman, 2017) and the US are endorsing the OECD AI Principles of transparency and explainability (OECD, 2019). Consequently, medical decision-supporting algorithms require verifying that this is not the result of exploiting data artifacts and that the high accuracy of classification decisions are explainable to avoid biased results (Lapuschkin et al., 2019, 2016). In the present work we used heat (or saliency) mapping, which is enabling perceptive interpretability to explain a classification result in terms of maps overlaid on the input (Tjoa and Guan, 2020). Regions in the input image contributing most to the classification result are highlighted in the heat map. From several methods currently available generating heat maps (Ribeiro et al., 2016; Simonyan et al., 2014; Springenberg et al., 2015; Zeiler and Fergus, 2014; Zintgraf et al., 2017), we based our proposed method on the deep Taylor decomposition (DTD) method (Montavon et al., 2017) which is a special case of layer-wise relevance propagation (LRP) (Bach et al., 2015). LRP, has a solid theoretical framework, has been extensively validated (Montavon, 2019; Samek et al., 2017) and can be efficiently implemented, enabling online heat map generation during training.

Besides indications from aforementioned studies, our experiments on Alzheimer's disease classification showed that CNNs might learn from (misleading) features outside the parenchyma or features introduced by the skull stripping algorithm. Thus, besides investigating how preprocessing steps including registration and skull stripping identify relevant features, we additionally present a novel relevance guided algorithm, mitigating the necessity and impact of skull stripping for classification of brain diseases. Based on its implementation this is referred to as Graz<sup>+</sup> technique (guided relevance by adaptive  $z^+$ -rule).

In summary, the specific contributions of this work are:

- CNN-based disease classification in a cohort of 128 patients with AD and 290 age-matched normal controls.
- Using subject-level 3D T1-based MR image data, differently preprocessed regarding registration and skull stripping.
- Graz<sup>+</sup> technique: A relevance-guided regularization technique for CNN classifiers to mitigate the impact of MRI preprocessing.
- Making the framework's source code freely available for reproducibility of the presented results.

## Methods

**Subjects.** Inclusion criteria for all participants was a diagnosis of probable or possible AD according to the NINCDS-ADRDA criteria (Knopman et al., 2001) and a complete MRI and study protocol as described in detail in (Damulina et al.,

2020). The healthy control (HC) group was selected from participants of a study in community-dwelling individuals. These volunteers were randomly selected from the community register, had a normal neurological status, and were without cerebrovascular attacks and dementia as previously described (Schmidt et al., 2003). This study was approved by the ethics committee of the Medical University of Graz (IRB00002556) and signed written informed consent was obtained from all study participants or their caregivers. The trial protocol for this prospective study was registered at the National Library of Medicine (trial identification number: NCT02752750).

**MR imaging.** Patients and controls were scanned using a consistent MRI protocol at 3 Tesla (Magnetom TimTrio; Syngo MR B17; Siemens Healthineers, Erlangen, Germany) using a 12-channel phased-array head coil. Structural imaging included a T1-weighted 3D MPRAGE sequence with 1 mm isotropic resolution (TR/TE/TI/FA = 1900 ms/2.19 ms/900 ms/9°, matrix = 176x224x256) and an axial FLAIR sequence (resolution of 1x1x3mm<sup>3</sup>) for the assessment of white matter abnormalities.

**Data selection.** Totally 132 patients with probable AD with 295 scans (Damulina et al., 2020) and 381 controls with 514 scans from an ongoing community dwelling study (Schmidt et al., 2003) were included in this retrospective study. From patients we excluded 12 MRIs because T1-weighted images were not available, and 14 scans because the image matrix was differently sized. Similarly, from controls we excluded 13 MRIs because of missing T1-weighted images, and 17 MRIs because of different image matrix sizes. Age-matching was achieved by excluding 5 scans of patients and 106 scans of controls, yielding 264 T1-weighted images from 128 patients with probable AD (mean age=71.9±8.5 years) and 378 MRIs from 290 healthy controls (mean age=71.3±6.4 years) for the subsequent deep learning analysis.

**Preprocessing.** Brain masks from T1-weighted MRIs were obtained using BET from FSL 6.0.3 with bias field/neck cleanup enabled and a fractional intensity threshold of 0.35 (Smith et al., 2004). T1-weighted images were registered to the MNI152 T1 template (**A**) affinely, using FSL flirt with 6 degrees of freedom and a correlation ratio based cost function, and (**B**) nonlinearly, using FSL fnirt with the *T1\_2\_MNI152\_2mm* configuration.

**Attention mask.** Our relevance-guided method is preconditioned by binary attention masks. We used entire brain masks obtained by FSL-BET to focus the classifiers to the intracranial volume.

**Classifier network.** We based our classifier on the 3D subject-level classifier network in (Wen et al., 2020). Although the proposed network is reported to perform quite well, the number of trainable parameters (42 millions) relative to the dataset size is high, thus rendering it prone to overfitting. Hence, the number and size of the convolutional

and fully connected layers were reduced until the network stopped overfitting on the training data and the validation accuracy started to drop. Batch normalization layers did not influence the performance of the network and were therefore removed. Finally, we replaced the max pooling layers by convolutional layers with striding as tested in (Springenberg et al., 2015). Avoiding max pooling layers improves the interpretability of networks (Montavon et al., 2018). Dropout was not applied in the network and all biases were constrained to be negative or zero. The final 3D classifier network is combining a single convolutional layer (kernel 3x3x3, 8 channels) with a down-convolutional layer (kernel 3x3x3, 8 channels, stridding 2x2x2) as the main building block. The overall network stacks 4 of these main building blocks followed by two fully connected layers (16 and 2 units) with totally 0.3 million trainable parameters. Each layer is followed by a Rectified Linear Unit (ReLU) nonlinearity, except for the output layer where a Softmax activation is applied.

**Heat mapping.** Heat maps were created based on the deep Taylor decomposition (DTD) method described in (Montavon et al., 2017). This method is equivalent to the layer-wise relevance propagation rule LRP- $\alpha_1\beta_0$  for networks like the one we used in this study. The principal idea of DTD is to compute a Taylor decomposition of the relevance at a given network layer onto the lower layer. The name “deep Taylor decomposition” comes from the iterative application of Taylor decomposition from the top layer down to the input layer (Montavon et al., 2018). The output of the Softmax layer of the classifier network defines the relevance that is redistributed with this saliency method. With DTD the relevance is routed only along the positively contributing parts of the network. This is a desired property because we want to focus the network on brain regions with features that cause the classification result. Nevertheless, it is of importance to select a heat mapping method that passes simple sanity checks and is dependent on the network and the training (Adebayo et al., 2020; Yona and Greenfeld, 2021). While lower layers could become less influencing on the saliency map (Sixt et al., 2020), DTD was shown to pass the sanity checks by computing saliency maps for all classes and then removing less relevant pixels from the final map (Gupta and Arora, 2019). Due to the nature of brain MRI data, we extended the currently available implementation of DTD from (Alber et al., 2019) to full 3D. The heat mapping method is used for both the relevance-guided classifier network and visualization.

**Relevance-guided classifier network.** The proposed relevance-guided network architecture focuses the classifier network on *relevant features* by extending the given network (cf. Figure 1 top) with a relevance map generator (cf. Figure 1 bottom). To this end we implemented the deep Taylor decomposition ( $z^+$ -rule) to generate the relevance maps of each input image depending on the classifier’s current parameters during training, yielding the Graz<sup>+</sup> technique (guided relevance by adaptive  $z^+$ -rule).

**Loss function.** In order to guide the training process by the attention mask ( $\mathbf{M}$ ), we extended the classifier’s categorical cross entropy loss ( $\text{loss}_{\text{CCE}}$ ) by a relevance-guided loss term to act as a regularizer:

$$\text{loss}_{\text{relevance}}(\mathbf{R}, \mathbf{M}) = -\mathbf{1}^T \text{vec}(\mathbf{R} \odot \mathbf{M}), \quad (1)$$

consequently yielding the total loss per data sample:

$$\begin{aligned} \text{loss}_{\text{Graz}^+} &= \text{loss}_{\text{relevance}} + \text{loss}_{\text{CCE}} \\ &= -\mathbf{1}^T \text{vec}(\mathbf{R} \odot \mathbf{M}) - \sum_{i=1}^{\text{outputs}} y_i \cdot \log(\hat{y}_i), \end{aligned} \quad (2)$$

where  $\mathbf{R}$  denotes the relevance heat map (3D shape),  $\mathbf{M}$  is the predefined binary attention mask obtained during image preprocessing (3D shape),  $\text{vec}(\mathbf{A})$  denotes the row major vector representation of  $\mathbf{A}$  resulting in a column vector (1D shape), and  $\mathbf{1}$  is a column vector where all elements are set to 1 (1D shape). The inner product of the transposed vector  $\mathbf{1}$  and the vector representation of  $\mathbf{R} \odot \mathbf{M}$  gives the scalar value  $\text{loss}_{\text{relevance}}$  (0D shape). The negative sign accounts for the maximization of the relevance inside the mask and  $\odot$  denotes the element-wise product. For the categorical cross entropy  $y_i$  is the target value of the  $i$ -th output class and  $\hat{y}_i$  its predicted value.


**Hyperparameter optimization and training.** Hyperparameter search on learning rate and learning rate schedule was done as proposed in CIFAR10-VGG11 experiments, in detail described in (Bouthillier et al., 2021). The batch size was omitted for consistent memory usage and exponential decay applied for learning rate schedule. Briefly, the hyperparameter optimization resulted in using the Adam optimizer with learning rate set to  $10^{-4}$ ,  $\gamma$  set to 1.0,  $\beta_1$  set to 0.9,  $\beta_2$  set to 0.999 and  $\epsilon$  set to  $10^{-7}$  (Kingma and Ba, 2015) for 60 epochs with a batch size of 8 for training in all configurations. Each model was end-to-end trained with standard loss minimization and error backpropagation. We trained models for 3 differently preprocessed T1-weighted input images

- in native subject space,
- linearly registered to MNI152 template and
- nonlinearly registered to MNI152 template

and all cases were tested in

- standard classifier network with native images,
- standard classifier network with the skull removed and
- our relevance-guided method with predefined attention masks,

creating overall nine models. No data augmentation was used.



**Fig. 1.** Schematic overview of the Graz<sup>+</sup> network and the adapted training process. A conventional classifier network (top) is extended by the heat map generator (bottom). For each classifier network layer a corresponding relevance redistribution layer with shared parameters and activations is attached to the generator network. The online calculated heat map is guiding the classifier training by adding a relevance sum inside the binary attention mask ( $\text{loss}_{\text{relevance}}$ ), which is added to the categorical cross entropy loss ( $\text{loss}_{\text{CCE}}$ ), yielding the total loss ( $\text{loss}_{\text{Graz}^+}$ ).  $\odot$  denotes the element-wise product.

**Cross validation.** AD and HC data were split up randomly into five folds without a separate test set, while maintaining all scans from one person in the same fold (Wen et al., 2020). Final folds were created by combining one fold from each cohort to ensure class distribution within. The difference in the class sizes was accounted for using a class weighting in the loss function.

**Model selection.** The optimal models based on the standard classifier networks were selected by highest validation classification accuracy. The relevance inside the attention mask threshold was set to 90% for the Graz<sup>+</sup> networks, enforcing models where most of the relevance is inside the intracranial volume.


**Relevance-weighted heat map representation.** Besides qualitatively investigating individual heat maps, we calculated mean heat maps and histogram for each mean heatmap. Starting with the bin with the highest relevance values, the bin contents were added up until 50% of all relevance was included. The lower value of the last bin added was used as the lower value for windowing the mean heatmap. All heat maps shown in this paper are overlaid on the MNI152 1mm

template and windowed to present the top 50% of relevance.

**Relevance density.** The relevance density describes the contribution of individual voxels of the heat map to the classification result. For all models we compare how many voxels are necessary to reach a certain level of explanation, e.g. how many voxels are needed to explain 85% of the total relevance.

**Volumetry.** For comparison between deep learning and logistic regression models for AD classification, we calculated whole brain, gray matter as well as ventricular volume using FSL-SIENAX with a fractional intensity threshold of 0.35 and bias field/neck cleanup enabled (Smith et al., 2002).

**Source code and data availability.** Source code for Graz<sup>+</sup> and the image preprocessing is available under [www.neuroimaging.at/explainable-ai](http://www.neuroimaging.at/explainable-ai). The MR images used in this paper are part of a clinical data set and therefore are not publicly available. Formal data sharing requests will be considered.



**Fig. 2.** Comparison of mean receiver operating characteristics curves for all nine configurations. The  $\text{Graz}^+$  models (blue) show higher values for the area under the curve (AUC in legend) compared to unmasked (purple) and masked (orange) configurations.

## Results

**Model performances.** Table 1 reports the mean performance for the cross validation setup of all tested configurations. In summary:

- While models with skull stripping perform better than those without, the  $\text{Graz}^+$  models yield even better balanced accuracy.
- The  $\text{Graz}^+$  model with linearly registered input had the highest balanced accuracy (86.19%), AUC (0.92) and also specificity (92.66%).
- Linear and nonlinear registration improves the balanced accuracy independently of skull stripping and utilization of  $\text{Graz}^+$ .
- The logistic regression model based on volumetric information for the entire brain, gray matter, and ventricular volume yielded a balanced accuracy of 82.00%, which is comparable or even outperforming some CNN models without skull stripping.

As the used dataset is nearly balanced (Saito and Rehmsmeier, 2015), the corresponding mean receiver operating characteristics (ROC) curves for these models are shown in Figure 2.

**Heat mapping.** Mean heat maps for classification decisions on cross validation holdout data sets for all trained models are shown in Figure 3, overlaid on the MNI152 1mm template. Individual heatmaps were nonlinearly transformed to the MNI152 space before averaging. Transformation information was obtained during T1-weighted image preprocessing. Visual inspection of the heat maps reveals that the preprocessing type (unmasked/masked/ $\text{Graz}^+$ ) yields substantially different results (columns), while the impact of the registration type (no registration/linear/nonlinear) is rather limited. Although mean heat maps in each column appear visually similar, applying registration to input MRIs improves the

balanced accuracy. When using the native T1-weighted images as input, the most relevant features are obtained in the scalp/skull outside brain parenchyma (unmasked configurations, left column). When skull stripping of the input MRIs is applied, the highest relevances are found in the cerebral and cerebellar cortex or generally adjacent to the brain-CSF-interface (middle column). While the aforementioned classifiers also show minor relevances in central brain regions, the maps from  $\text{Graz}^+$  show relevant regions exclusively within deep gray and white matter tissue adjacent to the ventricles (right column). Figure 4 shows multiple slices of mean heat maps for classification decisions of all cross validation holdout data sets for all trained models, overlaid on the MNI152 template.

**Relevance density.** Figure 5 shows that the  $\text{Graz}^+$  training increased the sparsity of the utilized features, where the 10% most relevant voxels (x-axis) explain approximately 20% (unmasked), 35% (masked) and 75% ( $\text{Graz}^+$ ) of the total relevance.

## Discussion

**Summary.** The present work investigated the mechanisms underlying brain disease classification by CNNs. Understanding the classifier's decision(s) is highly relevant, not only from an ethno-clinical but particularly from a legal perspective. We demonstrated how dramatic T1-weighted Alzheimer's disease classification is depending on volumetric features. Moreover, we show that preprocessing of neuroimaging data is decisive for feature identification because it introduces novel misleading features subsequently utilized for classification. The presented  $\text{Graz}^+$  technique is addressing these issues by focusing the feature identification on the intracranial space only. This yields higher classification accuracy than conventional CNN-methods, but more importantly, it substantially resolves the impact of MR image preprocessing.

### Impact to deep learning-based neuroimaging studies.

Our motivation for this work was driven by simple recurring questions in clinical brain MRI studies: Should the skull from a conventional T1-weighted MRI be stripped for further processing or should the entire MRI including skull and neck be used? Additionally, whether and which type of image registration is required or best as the next preprocessing step?

Showing that the preprocessing of MR images is crucial for the feature identification by CNNs has severe implications for neuroimaging based machine learning classifications. A majority of analysis pipelines apply skull stripping during image processing. This avoids the identification of features outside of the brain tissue, but in turn introduces new edges at the newly created brain mask, which might be subsequently used by the CNN for classification. We anticipate that decisions also might be misled by underlying contributors such as the implementation of the skull stripping algorithm, brain atrophy, but also might reflect visually not observable information as involuntary patients' movements. Generally, the


**Table 1.** Mean performance (in %) for the different models on all holdout data sets of cross validation. Highest values per column are highlighted in bold.

\*logistic regression by FSL-SIENAX (BET + tissue segmentation)


\*\*linear registration is applied during FSL-SIENAX processing to obtain scaling factor

AUC, area under the curve of the receiver operating characteristics.


Classifier	Skull stripping	Registration	Balanced accuracy	Sensitivity	Specificity	AUC
CNN	no	-	71.26±2.86%	55.55±7.51%	86.96±3.95%	0.75±0.02
		lin.	74.27±3.83%	63.13±9.05%	85.40±6.45%	0.80±0.05
		nonlin.	77.61±4.44%	64.79±5.02%	90.43±5.19%	0.85±0.06
CNN	yes	-	77.66±4.39%	69.70±7.65%	85.63±4.06%	0.83±0.05
		lin.	79.45±3.34%	76.87±4.81%	82.03±6.23%	0.86±0.05
		nonlin.	82.13±5.08%	73.47±7.89%	90.78±4.92%	0.88±0.05
CNN+Graz <sup>+</sup>	no	-	80.66±4.80%	74.95±7.85%	86.36±2.85%	0.88±0.04
		lin.	<b>86.19±6.01%</b>	79.73±10.72%	<b>92.66±3.73%</b>	<b>0.92±0.04</b>
		nonlin.	83.50±5.90%	77.16±8.95%	89.83±4.49%	0.90±0.04
Logistic Regression*	yes	lin.**	82.00±4.25%	<b>80.57±7.16%</b>	83.43±2.45%	0.90±0.04



**Fig. 3.** Mean heat maps (highest relevance in yellow, overlaid on MNI152 template) and balanced classification accuracy (percentage). Unmasked and masked CNN classifiers obtain relevant image features overwhelmingly from global volumetric information (left and center columns), whereas Graz<sup>+</sup> exclusively relies on deep gray and white matter tissue adjacent to the ventricles (right column). Heat maps are thresholded to the top 50% of the overall relevance. See Methods for description.



**Fig. 4.** MNI152 template overlaid by mean relevance maps (highest relevance in yellow) obtained for all nine models. Unmasked and masked MRI classifiers obtain relevant image features from volumetric information (left and center columns). In contrast, the proposed  $\text{Graz}^+$ -method bases the classifier's decision on deep brain image features, virtually independently of the registration method (right column). Heat maps are thresholded to the top 50% of the overall relevance. See Methods for description.



**Fig. 5.** The relevance density describes the contribution of individual voxels to the classification decision. Removal of scalp tissue voxels (orange) yields higher relevance density compared to unmasked T1 images (purple). The Graz<sup>+</sup>-models (blue) identify sparser but substantially more relevant voxels, which improves the classification accuracy.

source and extent of the newly introduced features remains unclear, however it was demonstrated that skull stripping algorithms can be biased by the patient cohort (Fennema-Notestine et al., 2006), thus, additionally biasing the classification.

Addressing these shortcomings, the proposed relevance-guided Graz<sup>+</sup> method identified regions of highest relevance in brain parenchyma while the balanced accuracy remained comparable or even better. Moreover, pooling data from rare diseases or generally small datasets often yield potentially spurious results and low replicability (Varoquaux, 2018). Its invariance from registration and skull stripping methods provides a usable method for CNN-based classification studies which might be practically useful when pooling data from different scanners and sites (Clarke et al., 2020) or assisting statistical harmonization (Dinsdale et al., 2021b; Pomponio et al., 2020).

**Neuroanatomical and Biophysical Interpretation.** This section highlights plausible mechanisms underlying CNN-based disease classification in AD by analyzing the neuroanatomical position of voxel relevance observed by heat mapping. The highest relevances were observed in the scalp for the CNN models using native (unmasked) input images. With skull stripping (masked), the most relevant voxels were found at the brain-CSF-interface, respectively, at the newly-introduced edges of the brain parenchyma. Anatomically, these regions are substantially overlapping with cortical gray matter, where atrophy is a well-known effect in AD. Cortical gray matter changes might be reflected in the masked CNNs decision, but seem rather implausible because of the small magnitude compared to global atrophy and ventricular enlargement. However, we cannot entirely rule out a secondary effect from the brain extraction algorithm biased by the patient cohort (Fennema-Notestine et al., 2006). Both CNN methods also identified some relevant voxel clusters in deep gray and white matter adjacent to the lateral ventricles

(center of the brain), which were substantially smaller.

Given the spatial distribution of the relevances, we argue that the two conventional CNN models are overwhelmingly sensitive for global volumetric features. Further evidence therefore comes from the complementary volumetric analysis using an established neuroimaging tool for brain segmentation (FSL-SIENAX) in a logistic regression model. The obtained balanced accuracy of 82% is on par with the top CNN results. Here the question arises whether these computational expensive CNNs just resample a refined volumetric measurement? The Graz<sup>+</sup>-based models identified regions with highest relevance mainly in deep gray and white matter located adjacent to the lateral ventricles. However, the anatomical/biophysical underpinnings of the decisions are less clear than in the conventional CNN models. Beside aforementioned contributions of volumetric features (AD progression is commonly paralleled by ventricular enlargement and global atrophy) also the T1-weighted contrast can pathologically change in AD (Besson et al., 1985). White matter hyperintensities (WMH) are commonly seen in brain MRI in older people and beside their underlying heterogeneous histopathology, they represent radiological correlates of cognitive and functional impairment (Prins and Scheltens, 2015). In a previous study, we found WMHs preferentially in a bilateral periventricular location, partly overlapping with the regions identified here by the Graz<sup>+</sup>-based models (Damulina et al., 2019). Furthermore, other plausible contributors are increased brain iron deposition in the deep gray matter (basal ganglia) of AD patients (Damulina et al., 2020) or cumulative gadolinium deposition of macrocyclic contrast agents (Kanda et al., 2014). Nevertheless, with the given setup we cannot definitely disentangle the underlying constituents and refer to the validation section below.

The relevance density analysis revealed that Graz<sup>+</sup>-based models learn much sparser features, subsequently needing less voxels for inferring classification decisions. Consequently, we hypothesize that the lack of misleading voxels from the scalp or newly-introduced edges is responsible for the increased accuracy.

**Related work.** With the availability of accessible large MRI databases from patients, such as the Alzheimer’s Disease Neuroimaging Initiative (ADNI), AIBL or OASIS databases, various studies using machine learning techniques exploiting structural imaging data have been published, formerly using *classical* machine learning classification methods (e.g. LDA, SVM) in combination with feature extraction methods based on tissue density (Klöppel et al., 2008), cortical surface (Eskildsen et al., 2013) and hippocampal measurements (Sørensen et al., 2016). Reported classification accuracies range between 75% and 100%, comprehensively summarized in (Rathore et al., 2017). Recently, interests switched to deep learning CNNs for (A) classification (Bäckström et al., 2018; Noor et al., 2020; Zhang et al., 2020), (B) classification with explanation (Böhle et al., 2019; Tang et al., 2019; Oh et al., 2019) and (C) regression with explanation (Dinsdale et al., 2021a) of AD. A recent review summarizes the state-of-the-art using CNNs for AD

classification, comparing various network architectures, input data and disease subtypes (Wen et al., 2020). Strictly in line with the data leakage analysis in this work we utilized stratified cross validation, while maintaining all datasets from one person in the same fold. Furthermore, we used the input MR images in their native spatial resolution, avoiding unpredictable influence from down- or resampling. While most of the analyzed studies are based on the ADNI dataset, our classification performance results are on par with both remaining 3D subject-level approaches without data leakage (Bäckström et al., 2018; Korolev et al., 2017).

The inconsistency between learned features with linear and nonlinear registration is systematically investigated in (Dinsdale et al., 2021a). They found that the use of nonlinearly registered images to train CNNs can drive the network by registration artifacts. However, the influence of further preprocessing steps on the resulting models and performances is less well known. Heat mapping using the LRP framework has been sparsely applied for explaining the underpinnings of an AD diagnosis in convolutional neural networks trained with structural MRI data beside the work of (Böhle et al., 2019). Heat maps obtained by two techniques (LRP and guided backpropagation) indicate relevant features in the parahippocampal gyrus but also adjacent to the brain-CSF interface, which is in line with our work.

Regularized heat map learning has been proposed before, however, differently to the Graz<sup>+</sup> method integrating a-priori knowledge with predefined attention masks. Technically, the gradient of the function learned by the network with respect to the current input can be interpreted as a heat map (Simonyan et al., 2014). Regularization of this input gradient was first introduced by (Drucker and Le Cun, 1992) as *double back-propagation*, which trains neural networks by not only minimizing the *energy* of the network but the rate of change of that energy with respect to the input features. In (Ross et al., 2017) this regularization was extended by selectively penalizing the gradient. Whereas (Sun et al., 2021) use LRP to create maps during training, which are multiplied with the corresponding input and then fed to the original classifier to dynamically find and emphasize important features. Furthermore, attention gated networks for medical image analysis have been proposed to automatically learn to focus on target structures of varying shapes and sizes (Schlemper et al., 2019).

**Validation.** Direct validation of the classifier’s decision is generally hardly feasible in the absence of a ground truth. While we anticipate a correspondence of the volumetric features with Alzheimer’s atrophy, this conclusion might not be final. However, in future work, indirect validation is possible using quantitative MRI parameters such as relaxometry, susceptibility, or magnetization transfer, where regional effects are known from ROI-based, voxel-based morphometry (VBM) or radiomics studies. While those methods statistically assess neuroanatomical features including ventricular enlargement or hippocampal atrophy, quantitative MRI pa-

rameters describe the underlying biophysical tissue composition. The effective relaxation rate  $R_2^*$  can assess increased iron deposition in the basal ganglia, a frequent finding in AD (Damulina et al., 2020). Consequently, the potential overlap with heat maps in those regions is better suited to disentangle biophysical tissue changes from atrophy. Optionally, direct validation of our method would require the generation of a cohort of realistic in silico phantoms (as recently used in the quantitative susceptibility mapping (QSM) image reconstruction challenge 2.0 (Marques et al., 2021)) with modulateable regional relaxation times in conjunction with an adjustable atrophy deformer (Khanal et al., 2017, 2016).

**Limitations.** Several aforementioned neuroimaging studies used the ADNI (or other publicly available) database for deep learning based classification. Generally, the clinical relevance of an automated AD classification is limited. The prodromal state of mild cognitive impairment (MCI) is preceding AD and identification of individuals rapidly progressing to AD (or differential diagnosis of frontotemporal dementia types) would be of higher importance for clinical management. We acknowledge the absence of an MCI group as a limitation and therefore provide the source code for the fast reproducibility using alternative network topologies, input data (quantitative MRI, PET), and other diseases. While aforementioned databases are designed multi-centrally, all MRI scans used in this paper were acquired with a single 3T scanner. Beside the underlying AD patient data, comparison with other studies is hampered by different network architectures, preprocessing and hyperparameter selection (Wen et al., 2020).

While this study only applied whole brain masks, more focused masks guiding the attention to e.g. the precuneus, the entorhinal cortex, the parietal lobe, the temporal lobe or the hippocampi are feasible, especially when regional a-priori knowledge for a certain pathology exists. Because of the explorative nature of the novel methodological framework we focused on the brain parenchyma. Organs outside the brain are more variable in size and shape, which render registration and ROI-definition more challenging. We originally developed Graz<sup>+</sup> for clinical brain studies, but its invariance to preprocessing might be even more pronounced beyond neuroimaging.

Lastly, the absence of CSF biomarkers or amyloid/Tau-PET for the AD diagnosis reduces the accuracy of the clinical diagnosis. However, AD diagnosis using the NINCDS-ADRDA criteria has a sensitivity of 81% and specificity of 70% as shown in clinico-pathological studies (Knopman et al., 2001).

## Conclusion

This work highlights that CNNs are not necessarily more efficient or better regarding classification accuracy than simple conventional volumetric features. However, the proposed relevance-guided approach is neutralizing the impact of MRI preprocessing from skull stripping and registration, rendering it a practically usable and robust method for CNN-based neu-

roimaging classification studies. Relevance-guiding focuses feature identification on the intracranial space only, yielding physiological plausible results and as a secondary effect the classification accuracy is higher.

## ACKNOWLEDGEMENTS

This study was funded by the Austrian Science Fund (FWF grant numbers: KLI523, P30134). This research was supported by NVIDIA GPU hardware grants.

## References

- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., Kim, B., 2020. Sanity Checks for Saliency Maps. arXiv:1810.03292 [cs, stat] URL: <http://arxiv.org/abs/1810.03292>. arXiv: 1810.03292.
- Alber, M., Lapuschkin, S., Seegerer, P., Hägele, M., Schütt, K.T., Montavon, G., Samek, W., Müller, K.R., Dähne, S., Kindermans, P.J., 2019. iNNvestigate Neural Networks! Journal of Machine Learning Research 20, 1–8. URL: <http://jmlr.org/papers/v20/18-540.html>.
- Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.R., Samek, W., 2015. On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation. PLoS One 10, e0130140. doi:10.1371/journal.pone.0130140.
- Besson, J.a.O., Corrigan, F.M., Foreman, E.I., Eastwood, L.M., Smith, F.W., Ashcroft, G.W., 1985. Nuclear Magnetic Resonance (NMR) II. Imaging in Dementia. The British Journal of Psychiatry 146, 31–35. URL: <https://www.cambridge.org/core/journals/the-british-journal-of-psychiatry/article/abs/nuclear-magnetic-resonance-nmr-ii-imaging-in-dementia/BFB2CEC043D56C97259797CC583FC34>, doi:10.1192/bjpp.146.1.31. publisher: Cambridge University Press.
- Biel, D., Brendel, M., Rubinski, A., Buerger, K., Janowitz, D., Dichgans, M., Franzmeier, N., Initiative (ADNI), ft.A.D.N., 2021. Tau-PET and in vivo Braak-staging as a prognostic marker in Alzheimer's disease. medRxiv , 2021.02.04.21250760URL: <https://www.medrxiv.org/content/10.1101/2021.02.04.21250760v1>, doi:10.1101/2021.02.04.21250760. publisher: Cold Spring Harbor Laboratory Press.
- Bouthillier, X., Delaunay, P., Bronzi, M., Trofimov, A., Nichyporuk, B., Szetó, J., Sepah, N., Raff, E., Madan, K., Voletti, V., Kahou, S.E., Michalski, V., Serdyuk, D., Arbel, T., Pal, C., Varoquaux, G., Vincent, P., 2021. Accounting for Variance in Machine Learning Benchmarks. arXiv:2103.03098 [cs, stat] URL: <http://arxiv.org/abs/2103.03098>. arXiv: 2103.03098.
- Braak, H., Alafuzoff, I., Arzberger, T., Kretschmar, H., Del Tredici, K., 2006. Staging of Alzheimer disease-associated neurofibrillary pathology using paraffin sections and immunocytochemistry. Acta Neuropathologica 112, 389–404. doi:10.1007/s00401-006-0127-z.
- Braak, H., Braak, E., 1991. Neuropathological staging of Alzheimer-related changes. Acta Neuropathologica 82, 239–259. doi:10.1007/BF00308809.
- Bäckström, K., Nazari, M., Gu, I.Y.H., Jakola, A.S., 2018. An efficient 3D deep convolutional network for Alzheimer's disease diagnosis using MR images, in: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018), pp. 149–153. doi:10.1109/ISBI.2018.8363543. ISSN: 1945-8452.
- Böhle, M., Eitel, F., Weygandt, M., Ritter, K., 2019. Layer-Wise Relevance Propagation for Explaining Deep Neural Network Decisions in MRI-Based Alzheimer's Disease Classification. Frontiers in Aging Neuroscience 11, 194. doi:10.3389/fnagi.2019.00194.
- Clarke, W.T., Mouglin, O., Driver, I.D., Rua, C., Morgan, A.T., Asghar, M., Clare, S., Francis, S., Wise, R.G., Rodgers, C.T., Carpenter, A., Muir, K., Bowtell, R., 2020. Multi-site harmonization of 7 tesla MRI neuroimaging protocols. NeuroImage 206, 116335. URL: <https://www.sciencedirect.com/science/article/pii/S1053811919309267>, doi:10.1016/j.neuroimage.2019.116335.
- Damulina, A., Pirpamer, L., Seiler, S., Benke, T., Dal-Bianco, P., Ransmayr, G., Struhal, W., Hofer, E., Langkammer, C., Duering, M., Fazekas, F., Schmidt, R., 2019. White Matter Hyperintensities in Alzheimer's Disease: A Lesion Probability Mapping Study. Journal of Alzheimer's disease: JAD 68, 789–796. doi:10.3233/JAD-180982.
- Damulina, A., Pirpamer, L., Soellradl, M., Sackl, M., Tinauer, C., Hofer, E., Enzinger, C., Gesierich, B., Duering, M., Ropele, S., Schmidt, R., Langkammer, C., 2020. Cross-sectional and Longitudinal Assessment of Brain Iron Level in Alzheimer Disease Using 3-T MRI. Radiology 296, 619–626. URL: <https://pubs.rsna.org/doi/10.1148/radiol.2020192541>, doi:10.1148/radiol.2020192541. publisher: Radiological Society of North America.
- Davatzikos, C., 2019. Machine learning in neuroimaging: Progress and challenges. NeuroImage 197, 652–656. URL: <https://www.sciencedirect.com/science/article/pii/S1053811918319621>, doi:10.1016/j.neuroimage.2018.10.003.
- Dinsdale, N.K., Bluemke, E., Smith, S.M., Arya, Z., Vidaurre, D., Jenkinson, M., Namburete, A.I.L., 2021a. Learning patterns of the ageing brain in MRI using deep convolutional networks. NeuroImage 224, 117401. doi:10.1016/j.neuroimage.2020.117401.
- Dinsdale, N.K., Jenkinson, M., Namburete, A.I.L., 2021b. Deep learning-based unlearning of dataset bias for MRI harmonisation and confound removal. NeuroImage 228, 117689. URL: <https://www.sciencedirect.com/science/article/pii/S1053811920311745>, doi:10.1016/j.neuroimage.2020.117689.
- Druker, H., Le Cun, Y., 1992. Improving generalization performance using double backpropagation. IEEE transactions on neural networks 3, 991–997. doi:10.1109/72.165600.
- Dubois, B., Villain, N., Frisoni, G.B., Rabinovici, G.D., Sabbagh, M., Cappa, S., Bejanin, A., Bombois, S., Epelbaum, S., Teichmann, M., Habert, M.O., Nordberg, A., Blennow, K., Galasko, D., Stern, Y., Rowe, C.C., Salloway, S., Schneider, L.S., Cummings, J.L., Feldman, H.H., 2021. Clinical diagnosis of Alzheimer's disease: recommendations of the International Working Group. The Lancet Neurology 20, 484–496. URL: [https://www.thelancet.com/journals/lanneur/article/PIIS1474-4422\(21\)00066-1/abstract](https://www.thelancet.com/journals/lanneur/article/PIIS1474-4422(21)00066-1/abstract), doi:10.1016/S1474-4422(21)00066-1. publisher: Elsevier.
- Eitel, F., Soehler, E., Bellmann-Strobl, J., Brandt, A.U., Ruprecht, K., Giess, R.M., Kuchling, J., Asseger, S., Weygandt, M., Haynes, J.D., Scheel, M., Paul, F., Ritter, K., 2019. Uncovering convolutional neural network decisions for diagnosing multiple sclerosis on conventional MRI using layer-wise relevance propagation. NeuroImage. Clinical 24, 102003. doi:10.1016/j.nicl.2019.102003.
- Eskildsen, S.F., Coupé, P., Garcia-Lorenzo, D., Fonov, V., Pruessner, J.C., Collins, D.L., Alzheimer's Disease Neuroimaging Initiative, 2013. Prediction of Alzheimer's disease in subjects with mild cognitive impairment from the ADNI cohort using patterns of cortical thinning. NeuroImage 65, 511–521. doi:10.1016/j.neuroimage.2012.09.058.
- Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S., 2017. Dermatologist-level classification of skin cancer with deep neural networks. Nature 542, 115–118. URL: <https://www.nature.com/articles/nature21056>, doi:10.1038/nature21056. number: 7639 Publisher: Nature Publishing Group.
- Fennema-Notestine, C., Ozyurt, I.B., Clark, C.P., Morris, S., Bischoff-Grethe, A., Bondi, M.W., Jerigan, T.L., Fischl, B., Segonne, F., Shattuck, D.W., Leahy, R.M., Rex, D.E., Toga, A.W., Zou, K.H., Brown, G.G., 2006. Quantitative evaluation of automated skull-stripping methods applied to contemporary and legacy images: effects of diagnosis, bias correction, and slice location. Human Brain Mapping 27, 99–113. doi:10.1002/hbm.20161.
- Goodman, B., Flaxman, S., 2017. European Union Regulations on Algorithmic Decision-Making and a "Right to Explanation". AI Magazine 38, 50–57. URL: <https://ojs.aaai.org/index.php/aimagazine/article/view/2741>, doi:10.1609/aimag.v38i3.2741. number: 3.
- Gupta, A., Arora, S., 2019. A Simple Saliency Method That Passes the Sanity Checks. arXiv:1905.12152 [cs, stat] URL: <http://arxiv.org/abs/1905.12152>. arXiv: 1905.12152.
- Hammernik, K., Klatzer, T., Kobler, E., Recht, M.P., Sodickson, D.K., Pock, T., Knoll, F., 2018. Learning a variational network for reconstruction of accelerated MRI data. Magnetic Resonance in Medicine 79, 3055–3071. doi:10.1002/mrm.26977.
- Henneman, W.J.P., Sluimer, J.D., Barnes, J., van der Flier, W.M., Sluimer, I.C., Fox, N.C., Scheltens, P., Vrenken, H., Barkhof, F., 2009. Hippocampal atrophy rates in Alzheimer disease: added value over whole brain volume measures. Neurology 72, 999–1007. doi:10.1212/01.wnl.0000344568.09360.31.
- Kanda, T., Ishii, K., Kawaguchi, H., Kitajima, K., Takenaka, D., 2014. High signal intensity in the dentate nucleus and globus pallidus on unenhanced T1-weighted MR images: relationship with increasing cumulative dose of a gadolinium-based contrast material. Radiology 270, 834–841. doi:10.1148/radiol.13131669.
- Karapinar Senturk, Z., 2020. Early diagnosis of Parkinson's disease using machine learning algorithms. Medical Hypotheses 138, 109603. doi:10.1016/j.mehy.2020.109603.
- Khanal, B., Ayache, N., Pennec, X., 2017. Simulating Longitudinal Brain MRIs with Known Volume Changes and Realistic Variations in Image Intensity. Frontiers in Neuroscience 11. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5360759/>, doi:10.3389/fnins.2017.00132.
- Khanal, B., Lorenzi, M., Ayache, N., Pennec, X., 2016. A biophysical model of brain deformation to simulate and analyze longitudinal MRIs of patients with Alzheimer's disease. NeuroImage 134, 35–52. doi:10.1016/j.neuroimage.2016.03.061.
- Kingma, D.P., Ba, J., 2015. Adam: A Method for Stochastic Optimization. ICLR .
- Kleesiek, J., Urban, G., Hubert, A., Schwarz, D., Maier-Hein, K., Bendszus, M., Biller, A., 2016. Deep MRI brain extraction: A 3D convolutional neural network for skull stripping. NeuroImage 129, 460–469. doi:10.1016/j.neuroimage.2016.01.024.
- Klöppel, S., Stnnington, C.M., Chu, C., Draganski, B., Scahill, R.I., Rohrer, J.D., Fox, N.C., Jack, C.R., Ashburner, J., Frackowiak, R.S.J., 2008. Automatic classification of MR scans in Alzheimer's disease. Brain: A Journal of Neurology 131, 681–689. doi:10.1093/brain/awm319.
- Knopman, D.S., DeKosky, S.T., Cummings, J.L., Chui, H., Corey-Bloom, J., Relkin, N., Small, G.W., Miller, B., Stevens, J.C., 2001. Practice parameter: diagnosis of dementia (an evidence-based review). Report of the Quality Standards Subcommittee of the American Academy of Neurology. Neurology 56, 1143–1153. doi:10.1212/wnl.56.9.1143.
- Korolev, S., Safiullin, A., Belyaev, M., Dodonova, Y., 2017. Residual and plain convolutional neural networks for 3D brain MRI classification, in: 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017), pp. 835–838. doi:10.1109/ISBI.2017.7950647. iSSN: 1945-8452.
- Lapuschkin, S., Binder, A., Montavon, G., Müller, K., Samek, W., 2016. Analyzing Classifiers: Fisher Vectors and Deep Neural Networks, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2912–2920. doi:10.1109/CVPR.2016.318. ISSN: 1063-6919.
- Lapuschkin, S., Wäldchen, S., Binder, A., Montavon, G., Samek, W., Müller, K.R., 2019. Unmasking Clever Hans predictors and assessing what machines really learn. Nature Communications 10, 1096. doi:10.1038/s41467-019-08987-4.
- Leung, K.K., Bartlett, J.W., Barnes, J., Manning, E.N., Ourselin, S., Fox, N.C., Alzheimer's Disease Neuroimaging Initiative, 2013. Cerebral atrophy in mild cognitive impairment and Alzheimer disease: rates and acceleration. Neurology 80, 648–654. doi:10.1212/WNL.0b013e318281ccd3.
- Marques, J.P., Meineke, J., Milovic, C., Bilgic, B., Chan, K.S., Hedouin, R., Zwaag, W.v.d., Langkammer, C., Schweser, F., 2021. QSM reconstruction challenge 2.0: A realistic in silico head phantom for MRI data simulation and evaluation of susceptibility mapping procedures. Magnetic Resonance in Medicine 86, 526–542. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28716>, doi:10.1002/mrm.28716. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.28716>.
- Montavon, G., 2019. Gradient-Based Vs. Propagation-Based Explanations: An Axiomatic Comparison, in: Samek, W., Montavon, G., Vedaldi, A., Hansen, L.K., Müller, K.R. (Eds.), Explainable AI: Interpreting, Explaining and Visualizing Deep Learning. Springer International Publishing, Cham. Lecture Notes in Computer Science, pp. 253–265. URL: [https://doi.org/10.1007/978-3-030-28954-6\\_13](https://doi.org/10.1007/978-3-030-28954-6_13), doi:10.1007/978-3-030-28954-6\_13.
- Montavon, G., Lapuschkin, S., Binder, A., Samek, W., Müller, K.R., 2017. Explaining nonlinear classification decisions with deep Taylor decomposition. Pattern Recognition

- 65, 211–222. URL: <https://www.sciencedirect.com/science/article/pii/S0031320316303582>, doi:10.1016/j.patcog.2016.11.008.
- Montavon, G., Samek, W., Müller, K.R., 2018. Methods for interpreting and understanding deep neural networks. *Digital Signal Processing* 73, 1–15. URL: <https://www.sciencedirect.com/science/article/pii/S1051200417302385>, doi:10.1016/j.dsp.2017.10.011.
- Noor, M.B.T., Zenia, N.Z., Kaiser, M.S., Mamun, S.A., Mahmud, M., 2020. Application of deep learning in detecting neurological disorders from magnetic resonance images: a survey on the detection of Alzheimer's disease, Parkinson's disease and schizophrenia. *Brain Informatics* 7, 11. URL: <https://doi.org/10.1186/s40708-020-00112-2>, doi:10.1186/s40708-020-00112-2.
- OECD, 2019. Artificial Intelligence in Society. OECD. URL: [https://www.oecd-ilibrary.org/science-and-technology/artificial-intelligence-in-society\\_eeefee77-en](https://www.oecd-ilibrary.org/science-and-technology/artificial-intelligence-in-society_eeefee77-en), doi:10.1787/eeefee77-en.
- Oh, K., Chung, Y.C., Kim, K.W., Kim, W.S., Oh, I.S., 2019. Classification and Visualization of Alzheimer's Disease using Volumetric Convolutional Neural Network and Transfer Learning. *Scientific Reports* 9, 18150. doi:10.1038/s41598-019-54548-6.
- Oldan, J.D., Jewells, V.L., Pieper, B., Wong, T.Z., 2021. Complete Evaluation of Dementia: PET and MRI Correlation and Diagnosis for the Neuroradiologist. *AJNR. American journal of neuroradiology* doi:10.3174/ajnr.A7079.
- Pomponio, R., Erus, G., Habes, M., Doshi, J., Srinivasan, D., Mamourian, E., Bashyam, V., Nasrallah, I.M., Satterthwaite, T.D., Fan, Y., Launer, L.J., Masters, C.L., Maruff, P., Zhuo, C., Völzke, H., Johnson, S.C., Frapp, J., Koutsouleris, N., Wolf, D.H., Gur, R., Gur, R., Morris, J., Albert, M.S., Grabe, H.J., Resnick, S.M., Bryan, R.N., Wolk, D.A., Shinohara, R.T., Shou, H., Davatzikos, C., 2020. Harmonization of large MRI datasets for the analysis of brain imaging patterns throughout the lifespan. *NeuroImage* 208, 116450. URL: <https://www.sciencedirect.com/science/article/pii/S1053811919310419>, doi:10.1016/j.neuroimage.2019.11.6450.
- Prins, N.D., Scheltens, P., 2015. White matter hyperintensities, cognitive impairment and dementia: an update. *Nature Reviews Neurology* 11, 157–165. URL: <https://www.nature.com/articles/nrneuro.2015.10>, doi:10.1038/nrneuro.2015.10.
- bandiera\_abtest: a Cg\_type: Nature Research Journals Number: 3 Primary\_atype: Reviews Publisher: Nature Publishing Group Subject\_term: Alzheimer's disease;Brain imaging;Dementia Subject\_term\_id: alzheimers-disease;brain-imaging;dementia.
- Rathore, S., Habes, M., Ifitkhar, M.A., Shacklett, A., Davatzikos, C., 2017. A review on neuroimaging-based classification studies and associated feature extraction methods for Alzheimer's disease and its prodromal stages. *NeuroImage* 155, 530–548. doi:10.1016/j.neuroimage.2017.03.057.
- Ribeiro, M.T., Singh, S., Guestrin, C., 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. *arXiv:1602.04938 [cs, stat]* URL: <http://arxiv.org/abs/1602.04938>. arXiv: 1602.04938.
- Ross, A.S., Hughes, M.C., Doshi-Velez, F., 2017. Right for the Right Reasons: Training Differentiable Models by Constraining their Explanations. *2662–2670* URL: <https://www.ijcai.org/proceedings/2017/371>.
- Saito, T., Rehmsmeier, M., 2015. The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets. *PLOS ONE* 10, e0118432. URL: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0118432>, doi:10.1371/journal.pone.0118432. publisher: Public Library of Science.
- Samek, W., Binder, A., Montavon, G., Lapuschkin, S., Müller, K., 2017. Evaluating the Visualization of What a Deep Neural Network Has Learned. *IEEE Transactions on Neural Networks and Learning Systems* 28, 2660–2673. doi:10.1109/TNNLS.2016.2599820. conference Name: IEEE Transactions on Neural Networks and Learning Systems.
- Scheltens, P., De Strooper, B., Kivipelto, M., Holstege, H., Chételat, G., Teunissen, C.E., Cummings, J., van der Flier, W.M., 2021. Alzheimer's disease. *Lancet (London, England)* 397, 1577–1590. doi:10.1016/S0140-6736(20)32205-4.
- Schlemper, J., Oktay, O., Schaap, M., Heinrich, M., Kainz, B., Glocker, B., Rueckert, D., 2019. Attention gated networks: Learning to leverage salient regions in medical images. *Medical Image Analysis* 53, 197–207. doi:10.1016/j.media.2019.01.012.
- Schmidt, R., Enzinger, C., Ropele, S., Schmidt, H., Fazekas, F., Austrian Stroke Prevention Study, 2003. Progression of cerebral white matter lesions: 6-year results of the Austrian Stroke Prevention Study. *Lancet (London, England)* 361, 2046–2048. doi:10.1016/S0140-6736(03)13616-1.
- Simonyan, K., Vedaldi, A., Zisserman, A., 2014. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. *ICLR* .
- Sixt, L., Granz, M., Landgraf, T., 2020. When Explanations Lie: Why Many Modified BP Attributions Fail. in: *Proceedings of the 37th International Conference on Machine Learning, PMLR*. pp. 9046–9057. URL: <https://proceedings.mlr.press/v119/sixt20a.html>. iSSN: 2640-3498.
- Sluimer, J.D., Vrenken, H., Blankenstein, M.A., Fox, N.C., Scheltens, P., Barkhof, F., van der Flier, W.M., 2008. Whole-brain atrophy rate in Alzheimer disease: identifying fast progressors. *Neurology* 70, 1836–1841. doi:10.1212/01.wnl.0000311446.61861.e3.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., Matthews, P.M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage* 23 Suppl 1, S208–219. doi:10.1016/j.neuroimage.2004.07.051.
- Smith, S.M., Zhang, Y., Jenkinson, M., Chen, J., Matthews, P.M., Federico, A., De Stefano, N., 2002. Accurate, robust, and automated longitudinal and cross-sectional brain change analysis. *NeuroImage* 17, 479–489. doi:10.1006/nimg.2002.1040.
- Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M., 2015. Striving for Simplicity: The All Convolutional Net. *arXiv:1412.6806 [cs]* URL: <http://arxiv.org/abs/1412.6806>. arXiv: 1412.6806.
- Sun, J., Lapuschkin, S., Samek, W., Zhao, Y., Cheung, N.M., Binder, A., 2021. Explanation-Guided Training for Cross-Domain Few-Shot Classification. in: *2020 25th International Conference on Pattern Recognition (ICPR)*, pp. 7609–7616. doi:10.1109/ICPR48806.2021.9412941. iSSN: 1051-4651.
- Sørensen, L., Igel, C., Liv Hansen, N., Osler, M., Lauritzen, M., Rostrup, E., Nielsen, M., Alzheimer's Disease Neuroimaging Initiative and the Australian Imaging Biomarkers and Lifestyle Flagship Study of Ageing, 2016. Early detection of Alzheimer's disease using MRI hippocampal texture. *Human Brain Mapping* 37, 1148–1161. doi:10.1002/hbm.23091.
- Tang, Z., Chuang, K.V., DeCarli, C., Jin, L.W., Beckett, L., Keiser, M.J., Dugger, B.N., 2019. Interpretable classification of Alzheimer's disease pathologies with a convolutional neural network pipeline. *Nature Communications* 10, 2173. URL: <https://www.nature.com/articles/s41467-019-10212-1>, doi:10.1038/s41467-019-10212-1. number: 1 Publisher: Nature Publishing Group.
- Tjoa, E., Guan, C., 2020. A Survey on Explainable Artificial Intelligence (XAI): Toward Medical XAI. *IEEE transactions on neural networks and learning systems* PP. doi:10.1109/TNNLS.2020.3027314.
- Varoquaux, G., 2018. Cross-validation failure: Small sample sizes lead to large error bars. *NeuroImage* 180, 68–77. URL: <https://www.sciencedirect.com/science/article/pii/S1053811917305311>, doi:10.1016/j.neuroimage.2017.06.061.
- Vieira, S., Pinaya, W.H.L., Mechelli, A., 2017. Using deep learning to investigate the neuroimaging correlates of psychiatric and neurological disorders: Methods and applications. *Neuroscience & Biobehavioral Reviews* 74, 58–75. URL: <https://www.sciencedirect.com/science/article/pii/S0149763416305176>, doi:10.1016/j.neubiorev.2017.01.002.
- Vogel, J.W., Young, A.L., Oxtoby, N.P., Smith, R., Ossenkopppe, R., Strandberg, O.T., La Joie, R., Aksam, L.M., Grothe, M.J., Iturria-Medina, Y., Alzheimer's Disease Neuroimaging Initiative, Pontecorvo, M.J., Devous, M.D., Rabinovici, G.D., Alexander, D.C., Lyoo, C.H., Evans, A.C., Hansson, O., 2021. Four distinct trajectories of tau deposition identified in Alzheimer's disease. *Nature Medicine* doi:10.1038/s41591-021-01309-6.
- Wen, J., Thibeau-Sutre, E., Diaz-Melo, M., Samper-González, J., Routier, A., Bottani, S., Dormont, D., Durrleman, S., Burgos, N., Colliot, O., Alzheimer's Disease Neuroimaging Initiative, Australian Imaging Biomarkers and Lifestyle flagship study of ageing, 2020. Convolutional neural networks for classification of Alzheimer's disease: Overview and reproducible evaluation. *Medical Image Analysis* 63, 101694. doi:10.1016/j.media.2020.101694.
- Yona, G., Greenfield, D., 2021. Revisiting Sanity Checks for Saliency Maps. *arXiv:2110.14297 [cs]* URL: <http://arxiv.org/abs/2110.14297>. arXiv: 2110.14297.
- Zeiler, M.D., Fergus, R., 2014. Visualizing and Understanding Convolutional Networks. in: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (Eds.), *Computer Vision – ECCV 2014*, Springer International Publishing, Cham. pp. 818–833. doi:10.1007/978-3-319-10590-1\_53.
- Zhang, L., Wang, M., Liu, M., Zhang, D., 2020. A Survey on Deep Learning for Neuroimaging-Based Brain Disorder Analysis. *Frontiers in Neuroscience* 14. URL: <https://www.frontiersin.org/articles/10.3389/fnins.2020.00779/full>, doi:10.3389/fnins.2020.00779. publisher: Frontiers.
- Zintgraf, L.M., Cohen, T.S., Adel, T., Welling, M., 2017. Visualizing Deep Neural Network Decisions: Prediction Difference Analysis. *arXiv:1702.04595 [cs]* URL: <http://arxiv.org/abs/1702.04595>. arXiv: 1702.04595.

## Appendix A

Table A.1 shows search space and default values for the hyperparameter optimizations for all configurations.

## Appendix B

Table B.1 shows performance for the different models on all holdout data sets of cross validation.

**Table A.1.** Search space and default values for the hyperparameter optimizations for all configurations.

<b>Hyperparameter</b>	<b>Default</b>	<b>Space</b>
learning rate	$10^{-4}$	$\log(10^{-5}, 10^{-3})$
$\gamma$ of learning rate schedule	0.997	$\text{lin}(0.96, 1.0)$
$\beta_1$	0.9	-
$\beta_2$	0.999	-
batch size	8	-

**Table B.1.** Performance (in %) for the different models on all holdout data sets of cross validation.

\*logistic regression by FSL-SIENAX (BET + tissue segmentation)

\*\*linear registration is applied during FSL-SIENAX processing to obtain scaling factor

AUC, area under the curve of the receiver operating characteristics.

Classifier	Skull stripping	Registration	Cross validation index	Balanced accuracy	Sensitivity	Specificity	AUC
CNN	no	-	1	73.28%	54.90%	91.67%	0.74
			2	67.02%	48.15%	85.90%	0.72
			3	70.19%	49.02%	91.36%	0.77
			4	70.41%	56.60%	84.21%	0.78
			5	75.39%	69.09%	81.69%	0.76
CNN	no	lin.	1	77.49%	64.71%	90.28%	0.76
			2	67.81%	57.41%	78.21%	0.75
			3	73.24%	58.82%	87.65%	0.84
			4	74.07%	54.72%	93.42%	0.80
			5	78.73%	80.00%	77.46%	0.87
CNN	no	nonlin.	1	83.62%	68.63%	98.61%	0.92
			2	71.94%	59.26%	84.62%	0.78
			3	75.93%	66.67%	85.19%	0.80
			4	74.64%	58.49%	90.79%	0.83
			5	81.93%	70.91%	92.96%	0.91
CNN	yes	-	1	80.15%	68.63%	91.67%	0.84
			2	70.66%	59.26%	82.05%	0.76
			3	78.36%	76.47%	80.25%	0.82
			4	75.50%	64.15%	86.84%	0.80
			5	83.66%	80.00%	87.32%	0.90
CNN	yes	lin.	1	78.35%	70.59%	86.11%	0.89
			2	77.14%	72.22%	82.05%	0.80
			3	76.98%	82.35%	71.60%	0.81
			4	78.81%	77.36%	80.26%	0.86
			5	85.98%	81.82%	90.14%	0.93
CNN	yes	nonlin.	1	89.09%	82.35%	95.83%	0.93
			2	74.07%	64.81%	83.33%	0.79
			3	79.99%	68.63%	91.36%	0.85
			4	81.99%	67.92%	96.05%	0.93
			5	85.48%	83.64%	87.32%	0.90
CNN+Graz <sup>+</sup>	no	-	1	86.32%	82.35%	90.28%	0.92
			2	79.06%	72.22%	85.90%	0.88
			3	76.03%	70.59%	81.48%	0.81
			4	75.50%	64.15%	86.84%	0.84
			5	86.39%	85.45%	87.32%	0.93
CNN+Graz <sup>+</sup>	no	lin.	1	92.44%	86.27%	98.61%	0.96
			2	89.67%	87.04%	92.31%	0.94
			3	75.82%	62.75%	88.89%	0.85
			4	83.22%	71.70%	94.74%	0.92
			5	89.82%	90.91%	88.73%	0.94
CNN+Graz <sup>+</sup>	no	nonlin.	1	89.79%	82.35%	97.22%	0.93
			2	79.13%	68.52%	89.74%	0.89
			3	76.54%	66.67%	86.42%	0.84
			4	80.78%	77.36%	84.21%	0.88
			5	91.23%	90.91%	91.55%	0.94
Log. Repr.*	yes	lin.**	1	80.31%	74.51%	86.11%	0.89
			2	76.50%	72.22%	80.77%	0.86
			3	79.34%	78.43%	80.25%	0.85
			4	86.16%	86.79%	85.53%	0.94
			5	87.71%	90.91%	84.51%	0.95