Predicting Body Mass Index From Structural MRI Brain Images Using a Deep Convolutional Neural Network

Pál Vakli[1]*, Regina J. Deák-Meszlényi[1], Tibor Auer[2] and Zoltán Vidnyánszky[1]

[1] *Brain Imaging Centre, Research Centre for Natural Sciences, Budapest, Hungary,* [2] *School of Psychology, Faculty of Health and Medical Sciences, University of Surrey, Guildford, United Kingdom*

In recent years, deep learning (DL) has become more widespread in the fields of cognitive and clinical neuroimaging. Using deep neural network models to process neuroimaging data is an efficient method to classify brain disorders and identify individuals who are at increased risk of age-related cognitive decline and neurodegenerative disease. Here we investigated, for the first time, whether structural brain imaging and DL can be used for predicting a physical trait that is of significant clinical relevance—the body mass index (BMI) of the individual. We show that individual BMI can be accurately predicted using a deep convolutional neural network (CNN) and a single structural magnetic resonance imaging (MRI) brain scan along with information about age and sex. Localization maps computed for the CNN highlighted several brain structures that strongly contributed to BMI prediction, including the caudate nucleus and the amygdala. Comparison to the results obtained via a standard automatic brain segmentation method revealed that the CNN-based visualization approach yielded complementary evidence regarding the relationship between brain structure and BMI. Taken together, our results imply that predicting BMI from structural brain scans using DL represents a promising approach to investigate the relationship between brain morphological variability and individual differences in body weight and provide a new scope for future investigations regarding the potential clinical utility of brain-predicted BMI.

Keywords: deep learning, convolutional neural networks, magnetic resonance imaging, body mass index, caudate nucleus, amygdala

INTRODUCTION

Over the last few years, the use of deep learning (DL) has become increasingly widespread in the analysis of neuroimaging data in several different application domains (Arbabshirani et al., 2017; Litjens et al., 2017; Shen et al., 2017; Zaharchuk et al., 2018; Davatzikos, 2019). DL is a branch of machine learning that allows the construction of computational models that learn to represent data at increasing levels of abstraction to solve specific tasks (LeCun et al., 2015; Goodfellow et al., 2016). Among DL methods, deep convolutional neural networks (CNNs) (LeCun et al., 1990;

Lecun et al., 1998), which are widely adopted in the computer vision community due to their capability to achieve outstanding object detection performance (Krizhevsky et al., 2012), represent a promising approach to analyzing brain imaging data in studies of psychiatric and neurological disorders (Vieira et al., 2017; Durstewitz et al., 2019). The majority of studies employing CNNs used structural and/or functional magnetic resonance imaging (MRI) data to examine patients with Alzheimer's disease and mild cognitive impairment (Gupta et al., 2013; Payan and Montana, 2015; Sarraf and Tofighi, 2016; Farooq et al., 2017; Meszlényi et al., 2017; Hosseini-Asl et al., 2018; Islam and Zhang, 2018; Basaia et al., 2019); although there are examples of studies classifying other mental disorders as well, such as attention-deficit hyperactivity disorder (Zou et al., 2017) and alcoholism (Wang et al., 2017).

The potential of these methods lies partly in that—in contrast to conventional mass univariate analytical methods—machine learning in general and DL in particular allow statistical inferences at the individual level (Vieira et al., 2017). Besides the diagnosis of brain disorders, machine learning can also be used to identify individual differences in the brain aging process (Cole and Franke, 2017; Cole et al., 2019). DL methods are increasingly prevalent in this application area as well, as CNNs can be used to predict the chronological age of individual subjects based on structural brain MRI scans with a mean absolute error (MAE) of 4.16 years (Cole et al., 2017). Comparable results can be obtained with CNNs using whole-brain functional connectivity patterns, derived from resting-state fMRI data, as input (Li et al., 2018; Vakli et al., 2018). These findings bear significance for two main reasons. First, they provide proof of concept that a single MRI scan contains information that is strongly related to chronological age (Cole and Franke, 2017). Second, they provide a means to quantify the individual risk of age-related cognitive decline and disease. In fact, several studies have shown that an increase in brain-predicted age relative to chronological age is associated with various neurological and psychiatric disorders, poorer physical fitness, and increased risk of mortality (Franke and Gaser, 2012; Koutsouleris et al., 2014; Cole et al., 2015, 2018; Habes et al., 2016; Löwe et al., 2016; Pardoe et al., 2017).

The above findings demonstrate how computational models aimed at predicting a certain biometric trait have potential clinical applicability. Here we investigated whether structural brain imaging and machine learning can be used for predicting a physical trait that is of significant clinical relevance—the body mass index (BMI) of the individual. The prevalence and disease burden of excessive body weight is on the rise globally (The GBD 2015 Obesity Collaborators, 2017), and there is extensive evidence showing a relationship between obesity—defined as a BMI greater than 30 kg/m^2—and brain health. In particular, a number of studies have shown that obesity and associated cardiovascular disease and metabolic disorders in midlife are related to cognitive impairment and dementia in later life (Pedditizi et al., 2016; Dye et al., 2017; Alford et al., 2018; Singh-Manoux et al., 2018). To date, a large number of studies using conventional neuroimaging methods have investigated the differences in brain structure and function between obese/overweight and lean individuals.

Increased BMI has been associated with reduced gray matter volume (Pannacciulli et al., 2006; Taki et al., 2008; Raji et al., 2010; Brooks et al., 2013) and white matter integrity (Stanek et al., 2011; Kullmann et al., 2015). Altered resting-state functional connectivity (Avery et al., 2017) and activation to visual food cues in brain regions involved in reward processing and inhibitory control (Carnell et al., 2012; Pursey et al., 2014; Val-Laillet et al., 2015) have also been described in obese individuals. A recent study has investigated the associations between obesity, regional gray matter volumes, and white matter microstructure, as assessed by MRI, in a large sample of 12,087 participants (Dekkers et al., 2019). The authors have found sex differences in the relationship between total body fat percentage and the volume of several subcortical regions of the brain reward system, and contrary to previous findings, a positive association between total body fat percentage and white matter microstructural coherence.

Training a machine learning algorithm to predict individual BMI based on brain imaging data has several potential applications. On the one hand, once sufficiently accurate prediction performance is achieved, it is possible to investigate which features (e.g., structural properties of the brain) contribute significantly to the predicted value. This has the potential to provide complementary information regarding the relationship between brain structure and body weight, besides conventional neuroimaging approaches. On the other hand, it can pave the way for potential clinical applications, inasmuch as the discrepancy between the true and the predicted BMI might be related to individual differences in food intake regulation and associated propensity for future weight gain. This would be analogous to that how the difference between brain-predicted and chronological age is used to quantify health risks.

Here we apply, for the first time to our knowledge, DL to predict individual BMI based on brain imaging data. In particular, we employ a CNN for BMI prediction based on T1-weighted structural MR images, as well as information about the participants' age and sex. This approach has the advantage of being able to use minimally preprocessed neuroimaging data as input and automatically learn a hierarchical set of representations suitable for solving the task at hand (LeCun et al., 2015), as opposed to conventional neuroimaging and machine learning methods that rely on *a priori* manual extraction of features from raw data (Vieira et al., 2017). Based on the findings discussed above, we hypothesized that BMI could be accurately predicted based on a single MRI bran scan, and hence a CNN can be trained to effectively perform this task on novel scans as well.

Once a well-performing model has been obtained and tested on new data, a logical next step is to try to make sense of why the model predicts what it predicts. While deep neural networks are usually regarded as "black boxes," it is possible to give reasonable explanations for their predictions without elucidating the underlying mechanisms (Lipton, 2016). Common approaches include projecting hidden layer activations back to input space to find patterns that excite feature maps the most (Zeiler and Fergus, 2014), examining the effect of occluding different parts of the input image on model performance (e.g., Vakli et al., 2018), or identifying those pixels in the input image that have the greatest impact on the model's predictions (e.g., Simonyan et al., 2013).

With regard to the latter approach, a particular method that has been used extensively in recent years to provide "visual explanations" for CNNs' decisions is Gradient-weighted Class Activation Mapping (Grad-CAM) (Selvaraju et al., 2017). This technique uses the gradient information flowing into the last convolutional layer of the CNN to highlight image regions that played an important role in predicting a certain target concept. Here we adapted this method to the context of regression based on 3D images to localize brain regions that made a significant contribution to BMI prediction.

Since the present study represents one of the first attempts to apply Grad-CAM for analyzing neuroimaging data, we also intended to investigate the neural underpinnings of individual differences in body weight using a more conventional neuroimaging approach and compare the obtained results. To this end, we performed automatic anatomical processing using the FreeSurfer software and general linear modeling to examine the relationship between brain morphology and BMI. FreeSurfer implements the automatic reconstruction of the cortical surface as well as subcortical structure segmentation using a probabilistic atlas (Dale et al., 1999; Fischl et al., 1999). The simultaneous application of the DL and automatic segmentation methods was motivated by the possibility that, as compared to this more conventional latter approach, using minimally preprocessed anatomical images and representation learning paired with gradient-based visualization would yield complementary evidence regarding the relationship between brain structure and body weight.

MATERIALS AND METHODS

Dataset

All analyses reported in this article include participants from the UK Biobank population cohort[1]. UK Biobank is a large prospective study comprising around 500,000 individuals recruited between 2006 and 2010 from across Great Britain who underwent physical and cognitive assessment, provided biological samples and completed questionnaires examining health and lifestyle (Allen et al., 2012). A subset of the participants (N = 22,392) underwent additional MRI from May 2014 until the data release in October 2018. Participants with a self-reported history of cancer, stroke, heart attack, deep-vein thrombosis, or pulmonary embolism diagnosed by a medical doctor (based on data-fields 2453, 6150, and 6152) were omitted from the current study. Additionally, only participants whose body mass indices were reported at the time of the imaging visit (data-field 21,001 instance 2) were included in the analyses. Finally, participants with a raw T1-weighted structural image deemed "unusable" by the UK Biobank team were also excluded. Image quality control on behalf of UK Biobank consisted of the rough manual review of T1 images supplemented by a beta-version automated quality control pipeline (Alfaro-Almagro et al., 2018). Eventually, 9518 females, aged between 45 and 80 years (mean $\pm$ SD = 62.11 $\pm$ 7.30 years), and 8420 males, aged

between 44 and 80 years (mean $\pm$ SD = 63.21 $\pm$ 7.59 years), were included in the present study. For females, BMI ranged between 13.39 and 58.70 kg/m^2 (mean $\pm$ SD = 26.15 $\pm$ 4.72 kg/m^2), while for males, it ranged between 16.67 and 58.04 kg/m^2 (mean $\pm$ SD = 27.03 $\pm$ 3.99 kg/m^2).

All participants provided informed consent to participate in the UK Biobank study. The UK Biobank Research Ethics Committee (REC) approval number is 11/NW/0382. Detailed information on the consent procedure of UK Biobank are available at the following URL: http://biobank.ctsu.ox.ac.uk/crystal/field.cgi?id=200.

Data Acquisition and Preprocessing
Neuroimaging

Data were acquired on Siemens Skyra 3T MRI scanners (Siemens Healthcare, Erlangen, Germany) at the UK Biobank imaging centers in Cheadle, Newcastle, and Reading. A standard Siemens 32-channel RF receive head coil was applied. The brain imaging protocol included a T1-weighted 3D magnetization-prepared rapid gradient echo (MPRAGE) sequence for structural imaging, using in-plane acceleration (iPAT = 2) and a field-of-view (FOV) of 208 $\times$ 256 $\times$ 256 with isotropic 1 mm spatial resolution.

Raw T1-weighted images were preprocessed by the UK Biobank team using an automated processing pipeline based on FSL tools (Jenkinson et al., 2012). The preprocessing pipeline included gradient distortion correction, cutting down the FOV, skull stripping, and non-linear transformation to MNI152 space (Alfaro-Almagro et al., 2018). In-house preprocessing was limited to reducing the size of the images to ease the computational burden of processing large 3D volumes. In particular, the "zoom" function of the multi-dimensional image processing package (scipy.ndimage) of the SciPy ecosystem[2] was used to resample each image by a factor of 0.5 using spline interpolation, resulting in images of shape 91 $\times$ 109 $\times$ 91 with isotropic 2 mm spatial resolution.

Body Mass Index

Data on weight were collected using a Tanita BC418MA body composition analyzer (Tanita Corporation of America, Inc., Arlington Heights, IL, United States). A Seca 240 cm height measure (Seca Deutschland, Hamburg, Germany) was used to obtain standing height measurement from participants. Body mass index was calculated as follows:

$$BMI = \text{weight in kilograms/height in meters}^2$$

Further details on the anthropometric measurements can be obtained from the following URL: http://biobank.ndph.ox.ac.uk/showcase/refer.cgi?id=146620.

Age and Sex

The age of each participant was derived from the date of birth (data-fields 34, 52) and the date of the imaging visit (data-field 21,003 instance 2) and was given in years with precision to the month. Sex was self-reported (data-field 31) and coded as 0 for female and 1 for male.

---

[1]https://www.ukbiobank.ac.uk/

[2]https://scipy.org/

Prediction of Body Mass Index
Neural Network Architecture

We used a CNN to predict BMI. The prediction of the model is based on three inputs from each subject:

1. T1-weighted brain image in MNI152 space, encoded in a Numpy[3] array of shape 91 × 109 × 91.
2. Chronological age of the participant in years with precision to the month.

[3]https://numpy.org/

3. Sex of the participant (0 for female and or 1 for male).

The output of the network is a single scalar corresponding to the predicted BMI of the subject.

A schematic illustration of the network architecture is given in **Figure 1**. The network comprises repeated blocks of 3D spatially separable convolutional layers followed by batch normalization (Ioffe and Szegedy, 2015) and rectified linear unit (ReLU) activation function (Nair and Hinton, 2010). In 3D spatially separable convolutional layers, instead of convolving the input
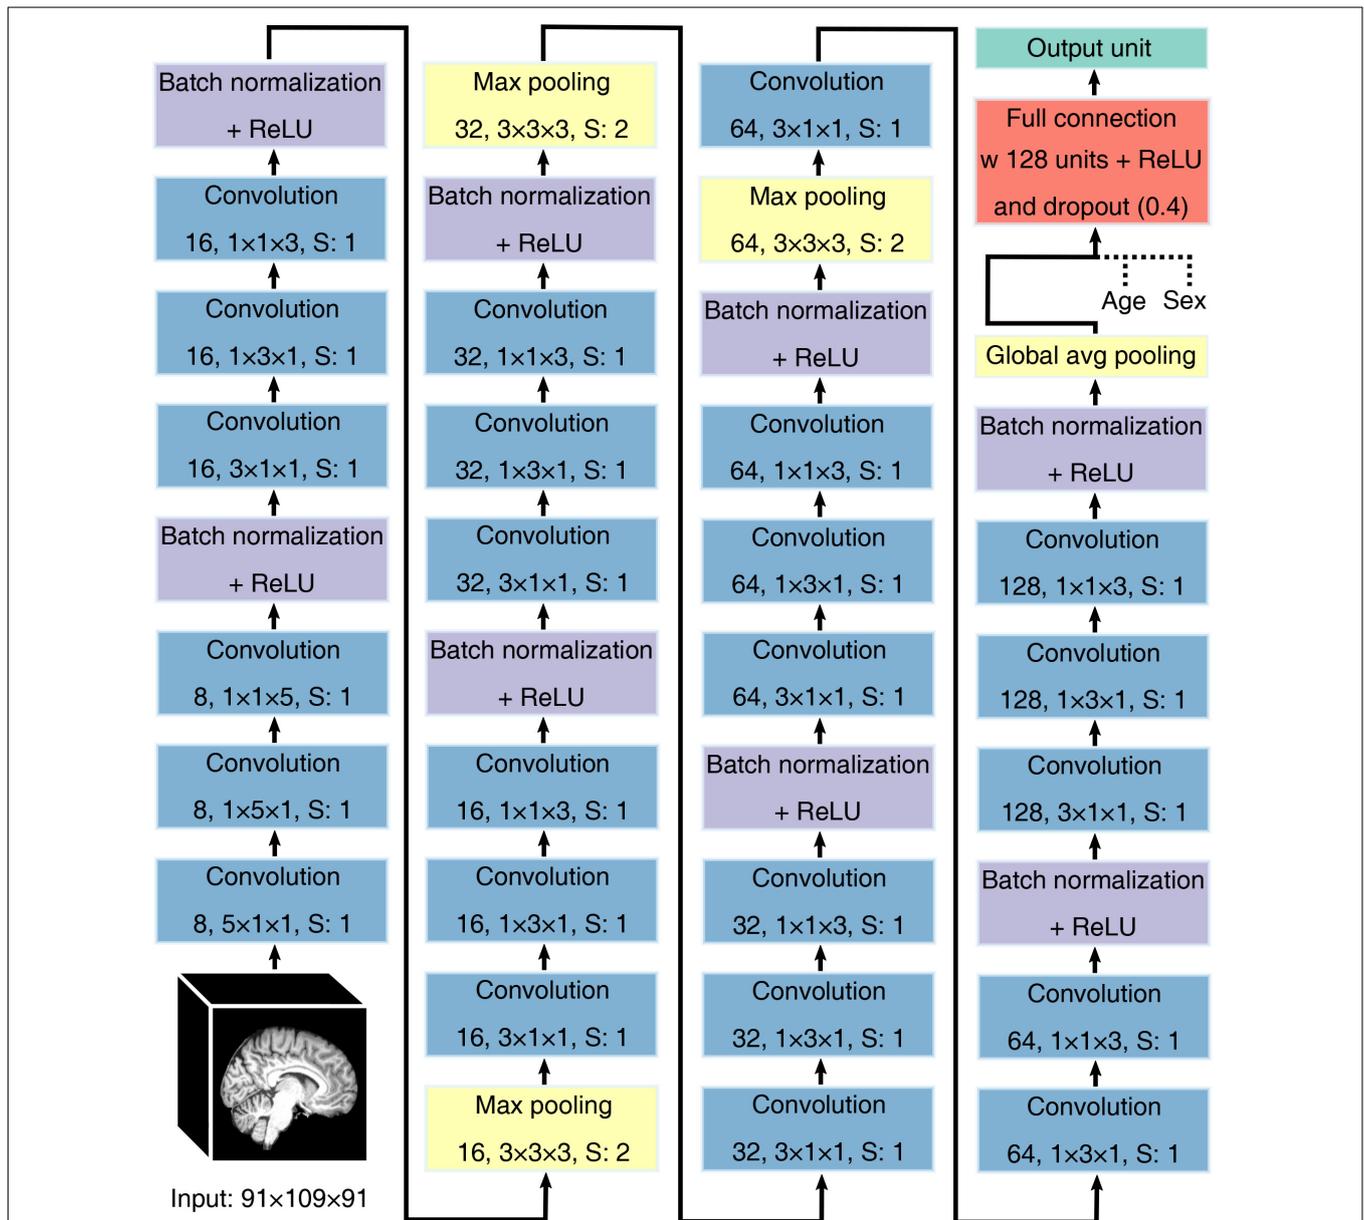

FIGURE 1 | Schematic illustration of the architecture of the convolutional neural network used for predicting body mass index. The network comprises repeated blocks of 3D spatially separable convolutional layers followed by batch normalization and ReLU, with every other block followed by a pooling layer to subsample the input. Global average pooling is used to map the feature maps of the last block to a vector (with a single scalar for each feature map) that is fed into a fully connected hidden layer followed by a single output unit for BMI prediction. Dashed lines denote concatenation, S denotes stride.

with filters of shape $N \times N \times N$, a cascade of three asymmetric filters of shapes $N \times 1 \times 1$, $1 \times N \times 1$, and $1 \times 1 \times N$ is used. Such a factorization of convolution operations reduces the computational cost by reducing the number of parameters (Szegedy et al., 2016) and has been used effectively in 3D medical image processing (Silva et al., 2018). Filter size is $N = 5$ (with a stride of 1) for the first set of convolution operations and $N = 3$ afterward. The number of filters is eight in the first convolutional layer and is doubled at regular intervals to enable the learning of a rich set of feature representations of the input brain image. All convolutional layers used SAME padding.

Every other batch normalization layer is followed by max pooling (filter shape $3 \times 3 \times 3$, stride = 2) to subsample the input images, and global average pooling is implemented after the last batch normalization layer to calculate the average intensity value of each feature map computed by the last convolutional layer. The output of this operation, along with the values representing age and sex, is fed into a fully connected hidden layer with 128 units and ReLU activation function. This hidden layer is connected to a single output unit, the activation of which corresponds to the predicted BMI value.

The CNN has 231,681 parameters overall, out of which 230,961 parameters are trainable. The model was implemented in Python using TensorFlow 1.13.[4] and the source code of the model along with the learnt parameters is available on GitHub: https://github.com/vaklip/cnn_3d_regression.

To examine whether information about age and sex was crucial for BMI prediction we also trained a network that was identical to the one described above, except that the values representing age and sex were not concatenated to the output of the global average pooling operation nor were they fed to the network in any other way.

Model Training

The weights of the convolutional and fully connected layers were initialized using Xavier initialization (Glorot and Bengio, 2010). The shifting and scaling parameters of the batch normalization layers were initialized to zeros and ones, respectively. The bias terms of the fully connected layers were initialized to 0.01. To train the network, we used mean squared error as the loss function, Adam optimizer (Kingma and Ba, 2014) with a learning rate of 0.0005 (momentum decay hyperparameter $\beta 1 = 0.9$, scaling decay hyperparameter $\beta 2 = 0.999$) and a batch size of eight. Dropout regularization (Wager et al., 2013; Srivastava et al., 2014) with a dropout rate of 0.4 was applied to the fully connected hidden layer during training.

The brain images of all participants were randomly assigned to disjoint training ($N = 13938$), validation ($N = 2000$), and test ($N = 2000$) sets. Only data in the training and validation sets were used for training and hyperparameter selection. The model was trained on the training set for a total of 50 epochs, and its performance was evaluated on the validation set after each epoch. A snapshot of the model parameters leading to the best validation set performance was restored and the final model was evaluated on the test set. Model performance is

characterized by the MAE, standard deviation of the absolute error (STDAE), coefficient of determination (R^2), root mean square error (RMSE), and Pearson's correlation coefficient (r) between the true and predicted BMI values.

A single NVIDIA Quadro M4000 GPU was used to train the CNN, with a runtime of about 1 h per epoch.

Transfer Learning

We used transfer learning to investigate the generalizability of our approach. Transfer learning refers to the method of training a neural network on one dataset (the source domain) and then adapting the model to a different dataset and/or task (the target domain) by transfer and fine-tuning of the previously learned model weights. In our case, the UK Biobank dataset constituted the source domain and the Information eXtraction from Images (IXI) dataset[5] including brain MR images from multiple sites in London constituted the target domain. We included the T1-weighted MR images of 269 subjects from the IXI dataset who fell into the age range corresponding to the UK Biobank sample: 177 females aged between 44 and 78 years (mean ± SD = 60.50 ± 8.32 years) and 115 males aged between 44 and 79 years (mean ± SD = 59.48 ± 9.05 years). These images were recorded using Philips Intera 3T ($N = 96$; Hammersmith Hospital) and Philips Gyroscan Intera 1.5T ($N = 173$; Guy's Hospital) scanners and a FOV of $150 \times 256 \times 256$ and spatial resolution of $1.2 \text{ mm} \times 0.938 \text{ mm} \times 0.938 \text{ mm}$. Images recorded at a third location (Institute of Psychiatry using a GE 1.5T system) were omitted from the current analysis due to the very low number of participants that matched the given age range ($N = 23$). In-house image preprocessing was limited to spatial normalization to MNI152 space and skull-stripping using the SPM12 toolbox[6] and custom-made scripts running on MATLAB 2015a (MathWorks Inc., Natick, MA, United States).

Images were randomly divided into disjoint training ($N = 197$), validation ($N = 36$), and test sets ($N = 36$). The weights of the network were initialized to those learnt on the UK Biobank dataset and then trained on the IXI dataset for 50 epochs, using data augmentation (random rotations of maximum 5 degrees and translations of 10 voxels). The neural network architecture and training hyperparameters were the same as those used for training on UK Biobank data. A snapshot of the model parameters leading to the best validation set performance (evaluated at the end of each epoch) was restored and the final model was evaluated on the test set.

Localizing Brain Regions Relevant for BMI Prediction

In order to obtain localization maps highlighting brain regions that are important for BMI prediction, we used a modified version of the Grad-CAM (Selvaraju et al., 2017). The Grad-CAM method aims to provide visual explanations for the decisions made by a wide variety of CNNs. It uses the gradients of a given target concept flowing into the final convolutional layer

---

[4]https://www.tensorflow.org

[5]https://brain-development.org/ixi-dataset/
[6]http://www.fil.ion.ucl.ac.uk/spm/software/spm12/

to produce a coarse localization map that highlights regions in the input image that are important for predicting that concept. We applied two modifications to the original method. First, we adapted it for processing 3D images, similarly to (Wang et al., 2019). We computed the gradient of the predicted BMI-score y with respect to the feature maps A^n of the last convolutional layer, and performed global average pooling on these gradients to obtain an importance weight α_n for each feature map:

$$\alpha_n = \frac{1}{Z} \sum_i \sum_j \sum_k \frac{\partial y}{\partial A_{ijk}^n} \qquad (1)$$

where Z is the number of units in a feature map. Then, the weighted combination of the features maps was calculated to obtain the localization map $L \in \mathbb{R}^{u \times v \times w}$:

$$L = \sum_n \alpha_n A^n \qquad (2)$$

In the original formulation of Grad-CAM, which was developed to provide class-discriminative visualizations, a ReLU was applied to L in order to highlight features that have a positive influence on the class of interest, as negative values would likely belong to other classes (Selvaraju et al., 2017). Here, since our CNN performed a regression task with a single output unit, and hence we were interested in features that have either positive or negative influence on predicted BMI, we omitted this step.

Localization maps were computed for each individual in the UK Biobank test set. They were upsampled to match the size of the input images using spline interpolation (for details, see section "Neuroimaging"). Intensity values were standardized to have zero mean and unit variance. As all brain images were registered to MNI152 space, a voxelwise grand average localization map across all test subjects could be computed. The resulting map was thresholded at two standard deviations from the mean and superimposed on the ch2bet MRIcron[7] template to visualize regions in the brain that made a strong contribution to BMI prediction. To investigate the robustness of the results, a grand average localization map was also computed for the training set. This localization map was visually indistinguishable from the one obtained for the test set.

Examining the Relationship Between BMI and Brain Volumetric and Morphometric Variability

Based on the visualization provided by the modified Grad-CAM method, we performed further exploratory analyses to investigate the association between BMI and morphological variability in the human brain using the UK Biobank data. To this end, we randomly selected a subset of 200 participants from the test set, with the only constraint being that the male–female ratio and the distribution of chronological age and BMI remain similar to those in the overall test set. We used FreeSurfer 6.0[8] to automatically parcellate the cortical surface and segment the

[7]https://people.cas.sc.edu/rorden/mricron/index.html
[8]http://surfer.nmr.mgh.harvard.edu

subcortical structures in the anatomical images of these subjects (Dale et al., 1999; Fischl et al., 1999). Then we investigated the relationship between different measures of cortical and subcortical anatomy—estimated by FreeSurfer—and the true BMI of participants, as detailed below.

Subcortical Segmentation
The volume-based stream of FreeSurfer (Fischl et al., 2002, 2004) was used to quantify the volumes of left and right hemisphere subcortical structures. Subcortical structures were selected for volumetric analysis based on the regions highlighted in the localization map produced by the modified Grad-CAM method. We computed partial correlations to examine the relationship between subcortical structure volume and BMI while controlling for chronological age, sex, and overall subcortical gray matter volume. We controlled for the former two variables since they were added as covariates to the CNN model which was therefore able to adjust for structural differences between individuals of different age and sex. Partial correlations were calculated using Statistica 13.4. (TIBCO Software Inc., Palo Alto, CA, United States).

Cortical Parcellation
The surface-based stream of FreeSurfer (Dale et al., 1999; Fischl et al., 1999) was used to construct models of the boundaries between white matter and cortical gray matter (the white surface), and between gray matter and the cerebrospinal fluid (the pial surface). The triangular tessellation of these surfaces allows for the calculation of several morphometric measures at each location (vertex) of the cortex, including cortical thickness, area, and curvature. We investigated the relationship between these three measures and BMI using FreeSurfer's Query, Design, Estimate, Contrast (QDEC) tool. Specifically, after smoothing individual subject data to the average surface with a 10-mm full-width at half maximum Gaussian kernel, a general linear model (GLM) with one of the morphometric measures as dependent variable was applied at each vertex, accounting for the effects of age, sex, and total cortical gray matter volume. False discovery rate (FDR) correction (threshold at 0.05) was applied to reduce Type I. errors associated with multiple comparisons.

Based on the grand average localization map, we directly investigated the association between the morphology of the right middle temporal gyrus and BMI. In particular, we computed partial correlations to examine the relationship between BMI and surface area, mean thickness and curvature while controlling for age, sex, and total cortical gray matter volume.
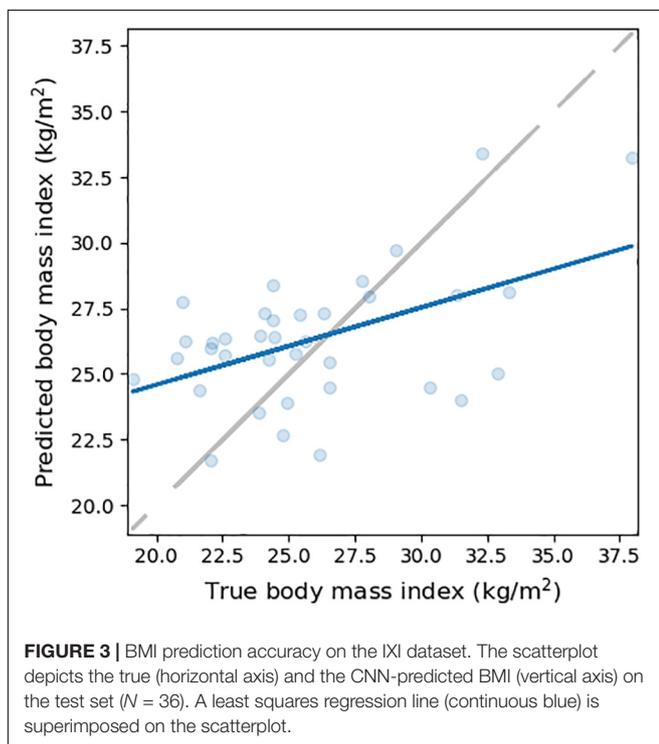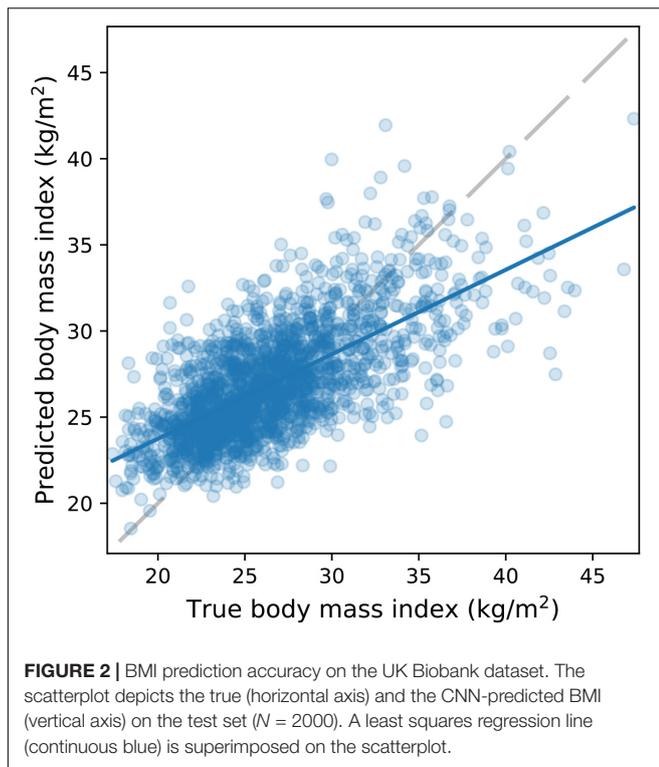
RESULTS

BMI Prediction
Overall, results showed that our CNN model can be used to predict BMI with high accuracy. Prediction error on the validation set reached a minimum after 32 epochs (MAE = 2.41 kg/m^2, STDAE = 1.93 kg/m^2). The model generalized well to the brain images in the test set

FIGURE 2 | BMI prediction accuracy on the UK Biobank dataset. The scatterplot depicts the true (horizontal axis) and the CNN-predicted BMI (vertical axis) on the test set (N = 2000). A least squares regression line (continuous blue) is superimposed on the scatterplot.

FIGURE 3 | BMI prediction accuracy on the IXI dataset. The scatterplot depicts the true (horizontal axis) and the CNN-predicted BMI (vertical axis) on the test set (N = 36). A least squares regression line (continuous blue) is superimposed on the scatterplot.

(**Figure 2**): MAE = 2.48 kg/m^2; STDAE = 2.09 kg/m^2; RMSE = 3.24 kg/m^2; Pearson r = 0.68; R^2 = 0.44.

When training the network without feeding information about age and sex to it, it took longer to reach a minimum

of prediction error on the validation set (after 41 epochs, MAE = 2.36 kg/m^2, STDAE = 2.09 kg/m^2). Nevertheless, the model generalized well to the test set images: MAE = 2.41 kg/m^2; STDAE = 2.11 kg/m^2; RMSE = 3.20 kg/m^2; Pearson r = 0.7; R^2 = 0.46.

When fine-tuning learned weights on the IXI dataset, validation error reached a minimum after 44 epochs (MAE = 2.53 kg/m^2; STDAE = 2.00 kg/m^2).We obtained reasonable BMI prediction on the IXI test set (**Figure 3**; MAE = 3.00 kg/m^2; STDAE = 2.12 kg/m^2; RMSE = 3.67 kg/m^2; Pearson r = 0.49; R^2 = 0.21), albeit it was below the performance obtained in the case of the UK Biobank dataset.

Localization Map

The grand average localization map across all the 2000 subjects' images in the test set is depicted in **Figure 4**. The map highlights several regions that, on average, have a strong influence on predicted BMI. These regions include the left caudate, the left medial temporal lobe in the vicinity of the amygdala, and the lateral surface of the right temporal cortex, encompassing the middle temporal gyrus.

Brain Volumetric and Morphometric Analyses

Based on the localization map, two subcortical regions, the left caudate and amygdala, were selected for volumetric analysis in a subset of the test subjects (**Figure 5**). On the one hand, there was no significant partial correlation between the volume of the caudate and the true BMI of the subjects when controlling for chronological age, sex, and overall subcortical gray matter volume (r = 0.028, p = 0.7). This may be accounted for by sex differences in the relationship between caudate volume and BMI (**Figure 5**, left panel). On the other hand, a significant partial correlation between the volume of the amygdala and BMI was observed (r = 0.19, p = 0.008), showing that increased BMI is associated with increased amygdalar volume.

Regarding the analysis of cortical morphometry, no significant association between BMI and cortical thickness or curvature was observed after correcting for multiple comparisons (FDR threshold at 0.05). However, a positive relationship was observed between BMI and the area of the isthmus cingulate in the right hemisphere (**Figure 6**). The direct tests (partial correlations) of the association between BMI and morphological measures of the right middle temporal gyrus yielded no significant results.

DISCUSSION

In this proof-of-concept study, we established that a deep CNN can be used to predict individual BMI with high accuracy, based on a single structural MRI brain scan and information about age and sex. This finding is in line with the results of several previous studies showing gray and white matter structural alterations in obese individuals (Brooks et al., 2013; Kullmann et al., 2015; Dekkers et al., 2019). We also demonstrated that gradient-based visualization can be used effectively to highlight brain regions that play an important role in BMI prediction.
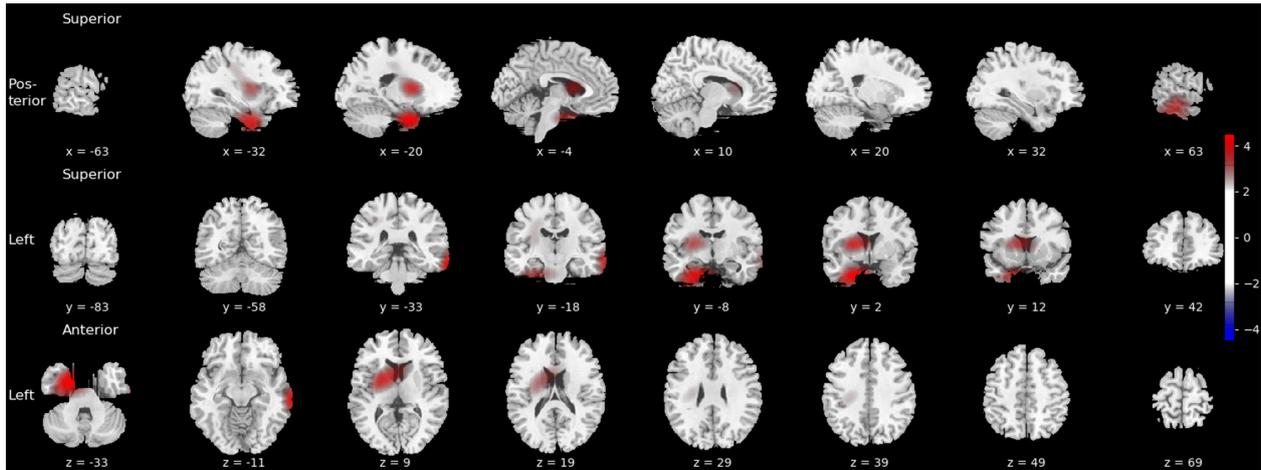
FIGURE 4 | Grand average localization map highlighting brain regions that strongly contribute to predicted BMI. Activation values are z-scored and thresholded at $|Z| > 2$. The localization map is superimposed on the ch2bet MRIcron template with MNI coordinates displayed below each slice.
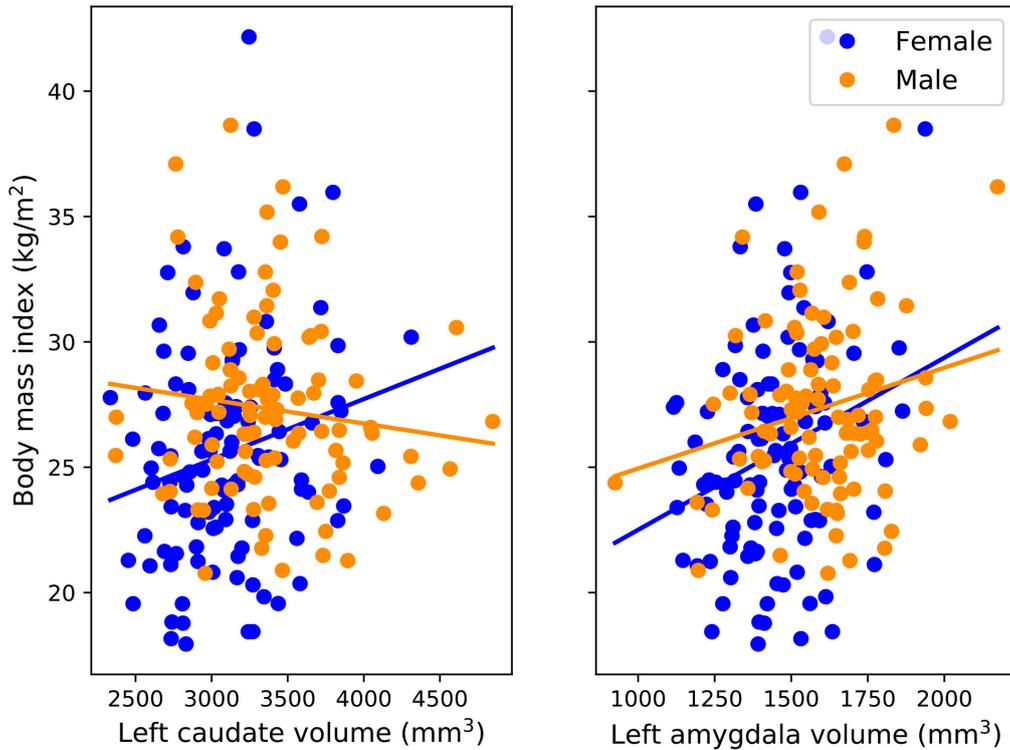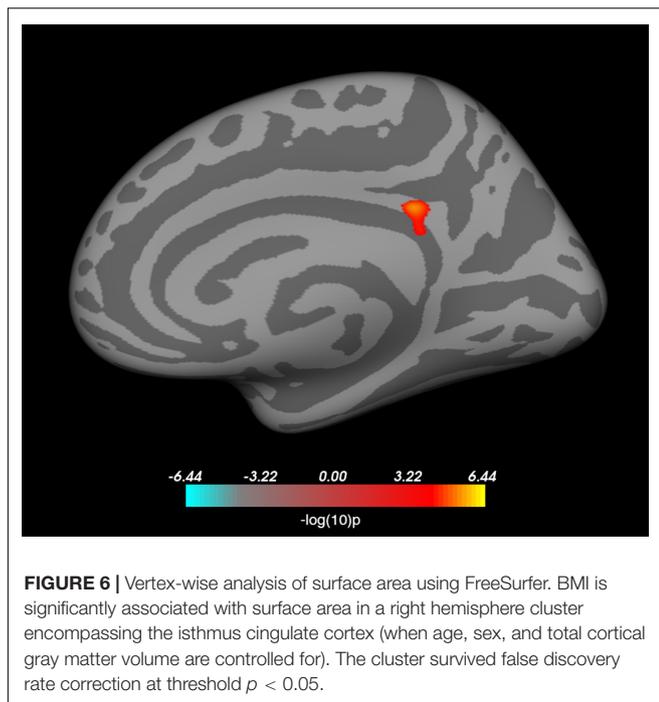

FIGURE 5 | BMI and subcortical volumes. Scatterplots depict the volumes of the caudate **(left panel)** and amygdala **(right panel)** in the left hemisphere and the true BMI values of male ($N = 93$) and female ($N = 107$) subjects in the test set.

More specifically, we used the Grad-CAM method, based on the gradient information flowing into the last convolutional layer of the CNN (Selvaraju et al., 2017), and adapted it to the context of regression using 3D images to identify brain regions that, on average, made a strong contribution to predicted BMI values. Our results suggest that, in addition to conventional neuroimaging methods and analytical techniques, the use of

DL along with visual explanations for model predictions is a suitable approach for identifying the brain structural correlates of individual variability in body weight.

In particular, the localization map produced by the Grad-CAM method highlighted a set of brain regions including a portion of the left medial temporal lobe in the vicinity of the amygdala. The relationship between amygdalar volume and

FIGURE 6 | Vertex-wise analysis of surface area using FreeSurfer. BMI is significantly associated with surface area in a right hemisphere cluster encompassing the isthmus cingulate cortex (when age, sex, and total cortical gray matter volume are controlled for). The cluster survived false discovery rate correction at threshold $p < 0.05$.

BMI was also confirmed by using FreeSurfer-based subcortical segmentation and partial correlation correcting for age and sex, which showed that higher BMI was associated with larger amygdalar volume. Previous studies using voxel-based (Taki et al., 2008) and tensor-based morphometry (Raji et al., 2010) found a relationship between BMI and the volume of gray and white matter in the medial temporal lobe. With regard to the amygdala, a positive relationship between BMI and amygdalar volume was already found in children and adolescents (Perlaki et al., 2018), young adults (Orsi et al., 2011), and elderly subjects (Widya et al., 2011); although a negative association has also been described (Kharabian Masouleh et al., 2016). Taken together, these results show that the DL approach paired with gradient-based visualization and more conventional neuroimaging methods provide converging evidence regarding the link between body weight and amygdalar structure. This is in accordance with the results of functional neuroimaging studies providing evidence for the involvement of the amygdala in processing visual food cues (van der Laan et al., 2011; Tang et al., 2012; van Bloemendaal et al., 2014).

Besides the commonalities, several discrepancies have been observed between the results of the Grad-CAM-based localization and the vertex-wise analysis using FreeSurfer. On the one hand, the vertex-wise analysis yielded a significant association between BMI and the surface area in a region corresponding to the isthmus cingulate in the right hemisphere. While at least one previous study reported a relationship between BMI and the morphology of the posterior cingulate cortex (Kharabian Masouleh et al., 2016), this region did not light up in the Grad-CAM-based localization map. On the other hand, several other brain structures were deemed important based on the localization map, in the case of which the conventional automatic brain segmentation approach failed to confirm an association with BMI, namely the lateral surface of the right temporal cortex and a region encompassing the left caudate nucleus. With regard to the latter, a previous study has shown that the volume of the caudate heads bilaterally show a positive association with BMI in men, after adjusting for age, lifetime alcohol intake, history of hypertension, and diabetes mellitus (Taki et al., 2008). Sex differences have also been shown to be manifest regarding the relationship between total body fat and caudate volume (Dekkers et al., 2019). Our results regarding the association with BMI are also indicative of such differences (**Figure 5**, left panel). In addition, the discrepancy between our observations with DL and conventional approaches is likely to stem from the differences in the applied methodologies as well. In our study, we used FreeSurfer for the automated segmentation of predefined subcortical structures and examined the linear relationship between BMI and a single scalar estimate of the volume of the caudate. FreeSurfer segmentation includes a series of pre-processing steps applied to the MRI volumes, followed by labeling the volumes based on a probabilistic atlas built from a set of hand-labeled images, as well as subject-specific measurements (Fischl et al., 2002, 2004). In contrast, the CNN is fed with minimally preprocessed images and learns a series of transformations to map those images to the corresponding BMI values. Each of these transformations map the representation of the input at one level into a representation at a slightly more abstract level (LeCun et al., 2015). Compared to the conventional automated brain segmentation methods, visualizations based on these more abstract representations may provide additional information with regard to the relationship between brain architecture and body weight. Similarly, several recent studies applied the Grad-CAM method to highlight brain regions that made an important contribution to predicting depression and epilepsy (Pominova et al., 2018), brain age (Bermudez et al., 2019), and Alzheimer's disease (Feng et al., 2018) based on structural MRI data.

Besides being a promising tool for neuroscientific investigation, brain-predicted BMI may also have practical utility. We managed to adapt the CNN model to a novel dataset, suggesting that our method is more generally applicable to a variety of different MR scanner types. Coming back to the relationship between the amygdala and body weight, this brain structure has been shown to be involved in the evaluation of food cues (Siep et al., 2009) and to constitute a part of a neural circuitry involved in the regulation of food craving (Dietrich et al., 2016). In a recent review, it has been argued that structures of the medial temporal lobe, in particular the amygdala and the hippocampus, may play an important role in the regulation of body weight, and that the amygdala is crucial for the regulation of feeding behavior based on environmental cues (Coppin, 2016). Based on the localization map produced by the Grad-CAM method, it is reasonable to hypothesize that brain-predicted BMI may be related to individual differences in the processing of food stimuli and cue-induced feeding. On this basis, one intriguing possibility is that increased brain-predicted BMI relative to the actual BMI might reflect a greater propensity to weight gain. This mode of application is similar to how the

difference between brain-predicted and chronological age might have clinical utility (Cole and Franke, 2017). However, it is important to note that brain structural alterations might not be the cause but the consequence of obesity. In fact, obesity-driven neuroinflammation has been shown to affect several brain regions including the hippocampus and the amygdala (Guillemot-Legris and Muccioli, 2017). Further research is necessary to examine whether and how brain-predicted BMI is related to pathophysiological processes and eating behavior.

CONCLUSION

Our findings provide proof of concept that individual BMI can be predicted with high accuracy from a single MRI scan using DL methods and suggest a relationship between the morphology of subcortical structures and body weight.

DATA AVAILABILITY STATEMENT

Publicly available datasets were analyzed in this study. This data can be found here: this research has been conducted using the UK Biobank Resource under Application Number 27236. All analyses reported in this paper include participants from the UK Biobank population cohort (https://www.ukbiobank.ac.uk/). The source code of the presented model along with the learnt parameters is available on GitHub: https://github.com/vaklip/cnn_3d_regression.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by UK Biobank Research Ethics Committee (REC; approval number: 11/NW/0382). The participants provided their written informed consent to participate in this study.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct and intellectual contribution to the work, and approved it for publication.

REFERENCES

Alfaro-Almagro, F., Jenkinson, M., Bangerter, N. K., Andersson, J. L. R., Griffanti, L., Douaud, G., et al. (2018). Image processing and quality control for the first 10,000 brain imaging datasets from UK Biobank. *Neuroimage* 166, 400–424. doi: 10.1016/j.neuroimage.2017.10.034

Alford, S., Patel, D., Perakakis, N., and Mantzoros, C. S. (2018). Obesity as a risk factor for Alzheimer's disease: weighing the evidence. *Obes. Rev.* 19, 269–280. doi: 10.1111/obr.12629

Allen, N., Sudlow, C., Downey, P., Peakman, T., Danesh, J., Elliott, P., et al. (2012). UK Biobank: current status and what it means for epidemiology. *Health Policy Technol.* 1, 123–126. doi: 10.1016/j.hlpt.2012.07.003

Arbabshirani, M. R., Plis, S., Sui, J., and Calhoun, V. D. (2017). Single subject prediction of brain disorders in neuroimaging: promises and pitfalls. *Neuroimage* 145, 137–165. doi: 10.1016/j.neuroimage.2016.02.079

Avery, J. A., Powell, J. N., Breslin, F. J., Lepping, R. J., Martin, L. E., Patrician, T. M., et al. (2017). Obesity is associated with altered mid-insula functional connectivity to limbic regions underlying appetitive responses to foods. *J. Psychopharmacol.* 31, 1475–1484. doi: 10.1177/0269881117728429

Basaia, S., Agosta, F., Wagner, L., Canu, E., Magnani, G., Santangelo, R., et al. (2019). Automated classification of Alzheimer's disease and mild cognitive impairment using a single MRI and deep neural networks. *Neuroimage Clin.* 21:101645. doi: 10.1016/j.nicl.2018.101645

Bermudez, C., Plassard, A. J., Chaganti, S., Huo, Y., Aboud, K. S., Cutting, L. E., et al. (2019). Anatomical context improves deep learning on the brain age estimation task. *Magn. Reson. Imaging* 62, 70–77. doi: 10.1016/j.mri.2019.06.018

Brooks, S. J., Benedict, C., Burgos, J., Kempton, M. J., Kullberg, J., Nordenskjöld, R., et al. (2013). Late-life obesity is associated with smaller global and regional gray matter volumes: a voxel-based morphometric study. *Int. J. Obes.* 37, 230–236. doi: 10.1038/ijo.2012.13

Carnell, S., Gibson, C., Benson, L., Ochner, C. N., and Geliebter, A. (2012). Neuroimaging and obesity: current knowledge and future directions. *Obes. Rev.* 13, 43–56. doi: 10.1111/j.1467-789X.2011.00927.x

Cole, J. H., and Franke, K. (2017). Predicting age using neuroimaging: innovative brain ageing biomarkers. *Trends Neurosci.* 40, 681–690. doi: 10.1016/j.tins.2017.10.001

Cole, J. H., Leech, R., and Sharp, D. J. (2015). Prediction of brain age suggests accelerated atrophy after traumatic brain injury. *Ann. Neurol.* 77, 571–581. doi: 10.1002/ana.24367

Cole, J. H., Marioni, R. E., Harris, S. E., and Deary, I. J. (2019). Brain age and other bodily 'ages': implications for neuropsychiatry. *Mol. Psychiatry* 24, 266–281. doi: 10.1038/s41380-018-0098-1

Cole, J. H., Poudel, R. P. K., Tsagkrasoulis, D., Caan, M. W. A., Steves, C., Spector, T. D., et al. (2017). Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. *Neuroimage* 163, 115–124. doi: 10.1016/j.neuroimage.2017.07.059

Cole, J. H., Ritchie, S. J., Bastin, M. E., Hernández, M. C. V., Maniega, S. M., Royle, N., et al. (2018). Brain age predicts mortality. *Mol. Psychiatry* 23, 1385–1392. doi: 10.1038/mp.2017.62

Coppin, G. (2016). The anterior medial temporal lobes: their role in food intake and body weight regulation. *Physiol. Behav.* 167, 60–70. doi: 10.1016/j.physbeh.2016.08.028

Dale, A. M., Fischl, B., and Sereno, M. I. (1999). Cortical surface-based analysis: I. Segmentation and surface reconstruction. *Neuroimage* 9, 179–194. doi: 10.1006/nimg.1998.0395

Davatzikos, C. (2019). Machine learning in neuroimaging: progress and challenges. *Neuroimage* 197, 652–656. doi: 10.1016/j.neuroimage.2018.10.003

Dekkers, I. A., Jansen, P. R., and Lamb, H. J. (2019). Obesity, brain volume, and white matter microstructure at MRI: a cross-sectional UK Biobank study. *Radiology* 291, 763–771. doi: 10.1148/radiol.2019181012

Dietrich, A., Hollmann, M., Mathar, D., Villringer, A., and Horstmann, A. (2016). Brain regulation of food craving: relationships with weight status and eating behavior. *Int. J. Obes.* 40, 982–989. doi: 10.1038/ijo.2016.28

Durstewitz, D., Koppe, G., and Meyer-Lindenberg, A. (2019). Deep neural networks in psychiatry. *Mol. Psychiatry* 24, 1583–1598. doi: 10.1038/s41380-019-0365-9

Dye, L., Boyle, N. B., Champ, C., and Lawton, C. (2017). The relationship between obesity and cognitive health and decline. *Proc. Nutr. Soc.* 76, 443–454. doi: 10.1017/S0029665117002014

Farooq, A., Anwar, S., Awais, M., and Rehman, S. (2017). "A deep CNN based multi-class classification of Alzheimer's disease using MRI," in *Proceedings of the 2017 IEEE International Conference on Imaging Systems and Techniques (IST)*, Beijing, 1–6. doi: 10.1109/IST.2017.8261460

Feng, X., Yang, J., Lipton, Z. C., Small, S. A., Provenzano, F. A., and Initiative, A. D. N. (2018). Deep learning on MRI affirms the prominence of the hippocampal formation in Alzheimer's disease classification. *bioRxiv* [Priprint]. doi: 10.1101/456277

Fischl, B., Salat, D. H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., et al. (2002). Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* 33, 341–355. doi: 10.1016/S0896-6273(02)00569-X

Fischl, B., Sereno, M. I., and Dale, A. M. (1999). Cortical surface-based analysis: II: inflation, flattening, and a surface-based coordinate system. *Neuroimage* 9, 195–207. doi: 10.1006/nimg.1998.0396

Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Ségonne, F., Salat, D. H., et al. (2004). Automatically parcellating the human cerebral cortex. *Cereb. Cortex* 14, 11–22. doi: 10.1093/cercor/bhg087

Franke, K., and Gaser, C. (2012). Longitudinal changes in individual BrainAGE in healthy aging, mild cognitive impairment, and Alzheimer's disease. *GeroPsych* 25, 235–245. doi: 10.1024/1662-9647/a000074

Glorot, X., and Bengio, Y. (2010). "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, eds Y. W. Teh, and M. Titterington (Sardinia: PMLR), 249–256.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. Cambridge MA: MIT press.

Guillemot-Legris, O., and Muccioli, G. G. (2017). Obesity-induced neuroinflammation: beyond the hypothalamus. *Trends Neurosci.* 40, 237–253. doi: 10.1016/j.tins.2017.02.005

Gupta, A., Ayhan, M. S., and Maida, A. S. (2013). "Natural image bases to represent neuroimaging data," in *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28 ICML'13, (JMLR.org)*, Atlanta.

Habes, M., Janowitz, D., Erus, G., Toledo, J. B., Resnick, S. M., Doshi, J., et al. (2016). Advanced brain aging: relationship with epidemiologic and genetic risk factors, and overlap with Alzheimer disease atrophy patterns. *Transl. Psychiatry* 6:e775. doi: 10.1038/tp.2016.39

Hosseini-Asl, E., Ghazal, M., Mahmoud, A., Aslantas, A., Shalaby, A. M., Casanova, M. F., et al. (2018). Alzheimer's disease diagnostics by a 3D deeply supervised adaptable convolutional network. *Front. Biosci. Landmark Ed.* 23:584–596. doi: 10.2741/4606

Ioffe, S., and Szegedy, C. (2015). Batch normalization: accelerating deep network training by reducing internal covariate shift. *arXiv* [Preprint]. Available at: http://arxiv.org/abs/1502.03167 [accessed July 29, 2019].

Islam, J., and Zhang, Y. (2018). Brain MRI analysis for Alzheimer's disease diagnosis using an ensemble system of deep convolutional neural networks. *Brain Inf.* 5:2. doi: 10.1186/s40708-018-0080-3

Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., and Smith, S. M. (2012). FSL. *Neuroimage* 62, 782–790. doi: 10.1016/j.neuroimage.2011.09.015

Kharabian Masouleh, S., Arélin, K., Horstmann, A., Lampe, L., Kipping, J. A., Luck, T., et al. (2016). Higher body mass index in older adults is associated with lower gray matter volume: implications for memory performance. *Neurobiol. Aging* 40, 1–10. doi: 10.1016/j.neurobiolaging.2015.12.020

Kingma, D. P., and Ba, J. (2014). Adam: a method for stochastic optimization. *arXiv* [Preprint]. Available at: http://arxiv.org/abs/1412.6980 [accessed January 22, 2018].

Koutsouleris, N., Davatzikos, C., Borgwardt, S., Gaser, C., Bottlender, R., Frodl, T., et al. (2014). Accelerated brain aging in schizophrenia and beyond: a neuroanatomical marker of psychiatric disorders. *Schizophr. Bull.* 40, 1140–1153. doi: 10.1093/schbul/sbt142

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). "ImageNet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems 25*, eds F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Red Hook, NY: Curran Associates, Inc.), 1097–1105.

Kullmann, S., Schweizer, F., Veit, R., Fritsche, A., and Preissl, H. (2015). Compromised white matter integrity in obesity. *Obes. Rev.* 16, 273–281. doi: 10.1111/obr.12248

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521:nature14539. doi: 10.1038/nature14539

LeCun, Y., Boser, B. E., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W. E., et al. (1990). "Handwritten digit recognition with a back-propagation network," in *Advances in Neural Information Processing Systems 2*, ed. D. S. Touretzky (Burlington, MA: Morgan-Kaufmann), 396–404.

Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 2278–2324. doi: 10.1109/5.726791

Li, H., Satterthwaite, T. D., and Fan, Y. (2018). "Brain age prediction based on resting-state functional connectivity patterns using convolutional neural networks," in *Proceedings of the 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)* (Washington, DC: IEEE), 101–104. doi: 10.1109/ISBI.2018.8363532

Lipton, Z. C. (2016). The mythos of model interpretability. *arXiv* [Preprint]. Available at: https://arxiv.org/abs/1606.03490v3 [accessed October 3, 2019].

Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., et al. (2017). A survey on deep learning in medical image analysis. *Med. Image Anal.* 42, 60–88. doi: 10.1016/j.media.2017.07.005

Löwe, L. C., Gaser, C., Franke, K., and The Alzheimer's Disease Neuroimaging Initiative, (2016). The effect of the APOE genotype on individual *BrainAGE* in normal aging, mild cognitive impairment, and Alzheimer's disease. *PLoS One* 11:e0157514. doi: 10.1371/journal.pone.0157514

Meszlényi, R. J., Buza, K., and Vidnyánszky, Z. (2017). Resting State fMRI functional connectivity-based classification using a convolutional neural network architecture. *Front. Neuroinformatics* 11:61. doi: 10.3389/fninf.2017.00061

Nair, V., and Hinton, G. E. (2010). "Rectified linear units improve restricted Boltzmann machines," in *Proceedings of the 27th International Conference on International Conference on Machine Learning ICML'10*, Haifa, 807–814.

Orsi, G., Perlaki, G., Kovacs, N., Aradi, M., Papp, Z., Karadi, K., et al. (2011). Body weight and the reward system: the volume of the right amygdala may be associated with body mass index in young overweight men. *Brain Imaging Behav.* 5, 149–157. doi: 10.1007/s11682-011-9119-2

Pannacciulli, N., Del Parigi, A., Chen, K., Le, D. S. N. T., Reiman, E. M., and Tataranni, P. A. (2006). Brain abnormalities in human obesity: a voxel-based morphometric study. *Neuroimage* 31, 1419–1425. doi: 10.1016/j.neuroimage.2006.01.047

Pardoe, H. R., Cole, J. H., Blackmon, K., Thesen, T., and Kuzniecky, R. (2017). Structural brain changes in medically refractory focal epilepsy resemble premature brain aging. *Epilepsy Res.* 133, 28–32. doi: 10.1016/j.eplepsyres.2017.03.007

Payan, A., and Montana, G. (2015). Predicting Alzheimer's disease: a neuroimaging study with 3D convolutional neural networks. *arXiv* [Preprint]. Available at: https://arxiv.org/abs/1502.02506v1 [accessed September 26, 2019].

Pedditizi, E., Peters, R., and Beckett, N. (2016). The risk of overweight/obesity in mid-life and late life for the development of dementia: a systematic review and meta-analysis of longitudinal studies. *Age Ageing* 45, 14–21. doi: 10.1093/ageing/afv151

Perlaki, G., Molnar, D., Smeets, P. A. M., Ahrens, W., Wolters, M., Eiben, G., et al. (2018). Volumetric gray matter measures of amygdala and accumbens in childhood overweight/obesity. *PLoS One* 13:e0205331. doi: 10.1371/journal.pone.0205331

Pominova, M., Artemov, A., Sharaev, M., Kondrateva, E., Bernstein, A., and Burnaev, E. (2018). "Voxelwise 3D convolutional and recurrent neural networks for epilepsy and depression diagnostics from structural and functional MRI data," in *Proceedings of the 2018 IEEE International Conference on Data Mining Workshops (ICDMW)*, Singapore, 299–307. doi: 10.1109/ICDMW.2018.00050

Pursey, K. M., Stanwell, P., Callister, R. J., Brain, K., Collins, C. E., and Burrows, T. L. (2014). Neural responses to visual food cues according to weight status:

a systematic review of functional magnetic resonance imaging studies. *Front. Nutr.* 1:7. doi: 10.3389/fnut.2014.00007

Raji, C. A., Ho, A. J., Parikshak, N. N., Becker, J. T., Lopez, O. L., Kuller, L. H., et al. (2010). Brain structure and obesity. *Hum. Brain Mapp.* 31, 353–364. doi: 10.1002/hbm.20870

Sarraf, S., and Tofighi, G. (2016). Classification of Alzheimer's disease using fMRI data and deep learning convolutional neural networks. *arXiv* [Preprint]. Available at: https://arxiv.org/abs/1603.08631v1 [accessed September 26, 2019].

Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D. (2017). Grad-CAM: visual explanations from deep networks via gradient-based localization. *arXiv* [Preprint]. Available at: http://openaccess.thecvf.com/content_iccv_2017/html/Selvaraju_Grad-CAM_Visual_Explanations_ICCV_2017_paper.html [accessed July 31, 2019].

Shen, D., Wu, G., and Suk, H.-I. (2017). Deep learning in medical image analysis. *Annu. Rev. Biomed. Eng.* 19, 221–248. doi: 10.1146/annurev-bioeng-071516-044442

Siep, N., Roefs, A., Roebroeck, A., Havermans, R., Bonte, M. L., and Jansen, A. (2009). Hunger is the best spice: an fMRI study of the effects of attention, hunger and calorie content on food reward processing in the amygdala and orbitofrontal cortex. *Behav. Brain Res.* 198, 149–158. doi: 10.1016/j.bbr.2008.10.035

Silva, J. F., Silva, J. M., Guerra, A., Matos, S., and Costa, C. (2018). "Ejection fraction classification in transthoracic echocardiography using a deep learning approach," in *Proceedings of the 2018 IEEE 31st International Symposium on Computer-Based Medical Systems (CBMS)*, Karlstad, 123–128. doi: 10.1109/CBMS.2018.00029

Simonyan, K., Vedaldi, A., and Zisserman, A. (2013). Deep inside convolutional networks: visualising image classification models and saliency maps. *arXiv* [Preprint]. Available at: https://arxiv.org/abs/1312.6034v2 [accessed October 4, 2019].

Singh-Manoux, A., Dugravot, A., Shipley, M., Brunner, E. J., Elbaz, A., Sabia, S., et al. (2018). Obesity trajectories and risk of dementia: 28 years of follow-up in the Whitehall II Study. *Alzheimers Dement.* 14, 178–186. doi: 10.1016/j.jalz.2017.06.2637

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15, 1929–1958.

Stanek, K. M., Grieve, S. M., Brickman, A. M., Korgaonkar, M. S., Paul, R. H., Cohen, R. A., et al. (2011). Obesity is associated with reduced white matter integrity in otherwise healthy adults*. *Obesity* 19, 500–504. doi: 10.1038/oby.2010.312

Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. (2016). *Rethinking the Inception Architecture for Computer Vision*. Available at: https://www.cv-foundation.org/openaccess/content_cvpr_2016/html/Szegedy_Rethinking_the_Inception_CVPR_2016_paper.html [accessed July 29, 2019].

Taki, Y., Kinomura, S., Sato, K., Inoue, K., Goto, R., Okada, K., et al. (2008). Relationship between body mass index and gray matter volume in 1,428 healthy individuals. *Obesity* 16, 119–124. doi: 10.1038/oby.2007.4

Tang, D. W., Fellows, L. K., Small, D. M., and Dagher, A. (2012). Food and drug cues activate similar brain regions: a meta-analysis of functional MRI studies. *Physiol. Behav.* 106, 317–324. doi: 10.1016/j.physbeh.2012.03.009

The GBD 2015 Obesity Collaborators, (2017). Health effects of overweight and obesity in 195 countries over 25 years. *N. Engl. J. Med.* 377, 13–27. doi: 10.1056/NEJMoa1614362

Vakli, P., Deák-Meszlényi, R. J., Hermann, P., and Vidnyánszky, Z. (2018). Transfer learning improves resting-state functional connectivity pattern analysis using

convolutional neural networks. *Gigascience* 7:giy130. doi: 10.1093/gigascience/giy130

Val-Laillet, D., Aarts, E., Weber, B., Ferrari, M., Quaresima, V., Stoeckel, L. E., et al. (2015). Neuroimaging and neuromodulation approaches to study eating behavior and prevent and treat eating disorders and obesity. *Neuroimage Clin.* 8, 1–31. doi: 10.1016/j.nicl.2015.03.016

van Bloemendaal, L., Jzerman, R. G. I, ten Kulve, J. S., Barkhof, F., Konrad, R. J., Drent, M. L., et al. (2014). GLP-1 receptor activation modulates appetite- and reward-related brain areas in humans. *Diabetes* 63, 4186–4196. doi: 10.2337/db14-0849

van der Laan, L. N., de Ridder, D. T. D., Viergever, M. A., and Smeets, P. A. M. (2011). The first taste is always with the eyes: a meta-analysis on the neural correlates of processing visual food cues. *Neuroimage* 55, 296–303. doi: 10.1016/j.neuroimage.2010.11.055

Vieira, S., Pinaya, W. H. L., and Mechelli, A. (2017). Using deep learning to investigate the neuroimaging correlates of psychiatric and neurological disorders: methods and applications. *Neurosci. Biobehav. Rev.* 74, 58–75. doi: 10.1016/j.neubiorev.2017.01.002

Wager, S., Wang, S., and Liang, P. S. (2013). "Dropout Training as Adaptive Regularization," in *Advances in Neural Information Processing Systems 26*, eds C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Red Hook, NY: Curran Associates, Inc), 351–359.

Wang, J., Knol, M., Tiulpin, A., Dubost, F., De, M. B., Vernooij, M., et al. (2019). Grey matter age prediction as a biomarker for risk of dementia: a population-based study. *bioRxiv* [Preprint]. doi: 10.1101/518506

Wang, S.-H., Lv, Y.-D., Sui, Y., Liu, S., Wang, S.-J., and Zhang, Y.-D. (2017). Alcoholism detection by data augmentation and convolutional neural network with stochastic pooling. *J. Med. Syst.* 42:2. doi: 10.1007/s10916-017-0845-x

Widya, R. L., de Roos, A., Trompet, S., de Craen, A. J., Westendorp, R. G., Smit, J. W., et al. (2011). Increased amygdalar and hippocampal volumes in elderly obese individuals with or at risk of cardiovascular disease. *Am. J. Clin. Nutr.* 93, 1190–1195. doi: 10.3945/ajcn.110.006304

Zaharchuk, G., Gong, E., Wintermark, M., Rubin, D., and Langlotz, C. P. (2018). Deep learning in neuroradiology. *Am. J. Neuroradiol.* 39, 1776–1784. doi: 10.3174/ajnr.A5543

Zeiler, M. D., and Fergus, R. (2014). "Visualizing and understanding convolutional networks," in *Computer Vision – ECCV 2014*, eds D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars (Cham: Springer), 818–833. doi: 10.1007/978-3-319-10590-1_53

Zou, L., Zheng, J., Miao, C., Mckeown, M. J., and Wang, Z. J. (2017). 3D CNN based automatic diagnosis of attention deficit hyperactivity disorder using functional and structural MRI. *IEEE Access* 5, 23626–23636. doi: 10.1109/ACCESS.2017.2762703