Published in final edited form as:

Proc IEEE Int Symp Biomed Imaging. 2020 April ; 2020: 420–423. doi:10.1109/isbi45749.2020.9098356.

BRAIN AGE ESTIMATION USING LSTM ON CHILDREN'S BRAIN MRI

Sheng He^{*}, Randy L. Gollub[†], Shawn N. Murphy[†], Juan David Perez^{*}, Sanjay Prabhu^{*}, Rudolph Pienaar^{*}, Richard L. Robertson^{*}, P. Ellen Grant^{*}, Yangming Ou^{*}

^{*} Boston Children's Hospital, Harvard Medical School, Boston, USA

[†] Massachusetts General Hospital, Harvard Medical School, Boston, USA

Abstract

Brain age prediction based on children's brain MRI is an important biomarker for brain health and brain development analysis. In this paper, we consider the 3D brain MRI volume as a sequence of 2D images and propose a new framework using the recurrent neural network for brain age estimation. The proposed method is named as 2D-ResNet18+Long short-term memory (LSTM), which consists of four parts: 2D ResNet18 for feature extraction on 2D images, a pooling layer for feature reduction over the sequences, an LSTM layer, and a final regression layer. We apply the proposed method on a public multisite NIH-PD dataset and evaluate generalization on a second multisite dataset, which shows that the proposed 2D-ResNet18+LSTM method provides better results than traditional 3D based neural network for brain age estimation.

Keywords

MRI; Age Prediction; ResNet; LSTM

1. INTRODUCTION

Predicting brain age from brain MRI is becoming an important biomarker for brain health and brain development analysis [1, 2]. The difference between the predicted and chronological age can be used to predict neurocognitive disorders [3] or brain anomaly caused by disease [2].

Inspired by deep learning, 3D-Convolutional Neural Network (3D-CNN) has been used to predict brain ages from a 3D brain MRI [4]. However, 3D-CNN requires extensive computations and more memory than 2D-CNN and it is hard to parallelize computations. A 2017 study used 3D-CNN on MRI to predict brain ages of N=2001 subjects aged from 18–90 years and reported an average prediction error of 4.16 years [2]. We are particularly interested in 0–20 years especially 0–6 years of age, where data is hard to get, the sample size is smaller, the brain is rapidly developing, and the inter-subject variability is bigger. Because of all these, 3D-CNN faces an additional risk for over-fitting.

To solve these problems, we propose a novel framework which considers the 3D brain MRI volume as a sequence of 2D images and uses recurrent neural network (such as the typical Long Short-Term Memory or LSTM [5]) to capture age information over the sequence of 2D slices. A similar approach (treating a 3D MRI as a series/sequence of 2D slices) was recently developed for voxel-wise abnormality detection [6], but we face a different task – patient-level prediction rather than voxel-wise classification, and hence a different formulation. In our framework, the feature representations of 2D slices are extracted by 2D-CNN which is easy to be parallelized and they are fed into LSTM to capture the global age information, which we assume is less likely to be trapped at local minima and less likely to over-fit in a small sample size in this age range.

In most brain age estimation studies, all available data are split into training and testing sets. Some of the training and testing subjects share imaging sites/protocols. So, although the exact testing subjects are unseen during training, the imaging sites/protocols may have already been exposed to training. In contrast, we use two multisite datasets: the NIH-Pediatric Data (NIH-PD) [7]¹ (GE/Siemens 1.5T scanner, 12 sites, spin-echo T1-weighted MRI), and the MGHBCCH dataset (Massachusetts General and Boston Children's Hospitals, Siemens 3.0T Trio scanners, T1-weighted MPRAGE sequence). The two datasets have very different imaging sites/protocols, and by training/testing on NIH-PD (the Discovery Cohort) and replicating on MGHBCCH data (the Replication Cohort, Fig 1), we fully evaluate 3D-CNN versus our proposed algorithm for the generalization ability in truly "unseen" dataset – not only testing subjects unseen but their imaging protocols/sites unseen.

2. METHODS

In this section, we describe the details of the proposed method, named as 2D-ResNet18+LSTM, which consists of four main parts (shown in Fig. 2): feature extraction by 2D residual network with 18 layers, a pooling for feature reduction, LSTM for context modeling over the sequence and a final regressor for age predication.

Feature extraction

Given a sequence of 2D slices of a 3D brain MRI volume, $\mathbf{X} = \{x_1, x_2, \dots, x_n\}$ (where n is the number of slices), 2D-ResNet [8] is used to extract feature representation. In this paper, we use ResNet18 without the fully-connected layer as the backbone to extract feature f_t for each 2D image x_t :

$$f_t = \text{ResNet}(\Theta, x_t) \quad (1)$$

where f_t is the feature representation of x_t and Θ parameters for ResNet18. We use ReLU after each convolutional layer as the activation function.

Pooling

Looking at 2D slices one by one at a constant pace may not be an optimal way to gather age information. It is likely that pooling information from several adjacent 2D slices can smooth

¹<https://www.nitrc.org/forum/forum.php?forumid=1152>

out noises in single slices and can better reveal true age information. This motivates us to combine features across adjacent 2D slices by a pooling operation. The most common pooling method is the average pooling, which is defined as:

$$p(m)_{avg} = \frac{1}{k} \sum_{j=k*m, \dots, k*m+k} f_j \quad (2)$$

where k is the kernel size of the pooling. Finally, the sequence is reduced to $P = \{p_1, p_2, \dots, p_m\}$ where $m = n/k$. In the experiments, we set $k = 3$.

LSTM

As mentioned above, each 3D brain MRI can be considered as a sequence of 2D slices, which, after pooling, are in m stacks of 2D slices represented by the feature vector $P = \{p_1, p_2, \dots, p_m\}$. This formulation motivates us to use recurrent neural networks, such as LSTM [5], to capture the age information. The LSTM cell (shown in Fig. 3) is defined as:

$$\begin{aligned} i_t &= \sigma(W_{ix} \cdot p_t + W_{ih} \cdot h_{t-1} + b_i) \\ f_t &= \sigma(W_{fx} \cdot p_t + W_{fh} \cdot h_{t-1} + b_f) \\ o_t &= \sigma(W_{ox} \cdot p_t + W_{oh} \cdot h_{t-1} + b_o) \\ g_t &= \varphi(W_{gx} \cdot p_t + W_{gh} \cdot h_{t-1} + b_g) \\ c_t &= f_t \odot c_{t-1} + i_t \odot g_t \\ h_t &= o_t \odot \varphi(c_t) \end{aligned} \quad (3)$$

where $\sigma(\cdot)$, $\varphi(\cdot)$ are the sigmoid and tangent $\tanh(\cdot)$ function, respectively. $\odot$ denotes an element-wise product. W_{*x} , W_{*h} and b_* are parameters which are learned during training. We use the Bidirectional LSTM with the hidden dimension size of 64, to balance between the performance and the computing cost.

Age regressor

All the outputs of LSTM cells are concatenated as the feature vector $\vec{H} = [h_1, h_2, \dots, h_t]$ and a fully connected layer is used as the regressor for age prediction.

Instance Normalization (IN)

Use 2D/3D batch normalization when the batch size is small (such as 1) may quickly explode the computer memory. Therefore, in this paper, we resort to instance normalization [9], which is defined as:

$$y = \frac{x - E[x]}{\sqrt{\text{Var}[x] + \epsilon}} \quad (4)$$

where the input x is a 2D/3D feature map. E and Var are the mean and variation, respectively. Instance Normalization is applied after each 2D/3D convolutional layer.

3. EXPERIMENTAL RESULTS

We use the NIH-PD dataset as the Discovery Cohort. We partitioned the 1212 subjects (0–22 years of age at the time of MRI scans) in the NIH-PD data into training (80%) and testing

sets (20%) to build the model and evaluate the model performance during the Discovery phase. Then, we use the MGHBCB dataset, which contains 428 normal-developing subjects 0–6 years of age, as a totally “unseen” Replication Cohort.

Our second setting to test the generalization ability of age prediction in the MGHBCB Replication Cohort (0–6 years of age) is to compare the accuracies by training on NIH-PD subjects 0–22 years and by training on NIH-PD subjects 0–6 years of age. The age distribution of each dataset is shown in Fig. 1.

T1-weighted MPRAGE MRI of every subject is registered to the SRI atlas (image size [128,160,112] and voxel size $1 \times 1 \times 1 \text{ mm}$ [10]). We compared using the SRI atlas (adults) and a 2-year-old pediatric atlas and found little differences in the final accuracies. We resize 2D images on the sequence to 50×50 voxels to reduce the computation time and the use of memory. The voxel values of each subject are normalized by subtracting the mean value of the subject’s image and dividing by the variance.

We use the corresponding 3D-ResNet18 as the baseline for comparison, which consists of 3D-convolutional layers with kernel size of $3 \times 3 \times 3$, global average pooling and a fully connected layer for age prediction. All networks, including 2D-ResNet18+LSTM and 3D-ResNet18 are trained with Mean Absolute Error (MAE) loss. The Adam optimization method is used with an initial learning rate of 0.0001, decreasing by half at every 15 epochs. The network is trained with 60 epochs. The batch size is set to 1.

The model is evaluated using two measurement metrics: Mean Absolute Error (MAE) and Cumulative Score (CS):

$$\begin{aligned} MAE &= \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i| \\ CS(\alpha) &= \frac{1}{N} \sum_{i=1}^N \{ |y_i - \hat{y}_i| \leq \alpha \} \times 100\% \end{aligned} \quad (5)$$

where y_i is the ground-truth age, $\hat{y}_i$ is the estimated age, α is the error level and N is the number of test samples. The $CS(\alpha)$ score is the percentage of subjects whose errors of age prediction are smaller than or equal to α . A smaller MAE and a higher $CS(\alpha)$ indicate a more accuracy algorithm.

3.1. Performance on the NIH-PD testing set (Discovery)

Table 1 shows the MAEs of different methods in different age groups on the NIH-PD testing set. The proposed 2D-ResNet18+LSTM provides lower MAEs in 6–10, 11–15 and 16–22 years, and gives a lower average MAEs than 3D-ResNet18. Fig. 4 shows the CS scores as a function of error levels α . The proposed 2D-ResNet18+LSTM provides better results than 3D-ResNet18 in terms of MAE and CS.

3.2. Generalization/Replication to the MGHBCB dataset

In this section, we provide the results of age estimation on the MGHBCB dataset using the model trained on the NIH-PD training set. This is to evaluate the generalization of 2D-ResNet18+LSTM and 3D-ResNet18 methods.

Table 2 shows the MAEs of these two methods on the MGHBCB dataset (0–6 years) after training on NIH-PD subjects 0–22 year of age. The proposed 2D-ResNet18+LSTM provides much better results than 3D-ResNet18 on the MGHBCB dataset. Fig. 5 shows the CS curves of these two methods and the proposed 2D-ResNet18+LSTM provides better results than 3D-ResNet18 with all different error levels α .

We repeat the above experiment in the Discovery (MGHBCB) Cohort, but using NIH-PD subjects 0–6 years for training. Table 3 shows the results. When the Discovery and Replicate Cohorts share the same age range, the average MAEs of 2D-ResNet18+LSTM and 3D-ResNet18 are approximately the same, but the computation time of 2D-ResNet18+LSTM is shorter than 3D-ResNet18 (3 hours versus 7 hours on an NVIDIA K80 GPU).

4. CONCLUSION

Recent age prediction is mainly on 18–90 years of age and using 2000–3000 brain MRIs. We focus on 0–22 especially 0–6 years where the sample size is smaller. In such circumstances, we shown that the proposed 2D-ResNet18+LSTM method provides comparative or better results than 3D-ResNet18.

While existing age prediction studies split the Discovery Cohort into training and testing, we further test the generalization ability in a truly unseen Replication Cohort (subjects unseen, imaging sites/protocols unseen). We show that the proposed 2D-ResNet18+LSTM method generalizes better to the Replication dataset when the age ranges of Discovery and Replication Cohorts are different (Table 2). This opens a window to train on a large-scale Discovery Cohort with a wider age range (such as 0–20 years) and to adapt to a Replication Cohort where the age range of interest is narrower (0–2 or 0–6 years) and the sample size is smaller.

When the age ranges of Discovery and Replication Cohorts are the same, 2D-ResNet18+LSTM has a faster computation, and offers a higher accuracy in both ends of the age range (Table 3) while 3D-ResNet18 has a higher accuracy in the middle of the age range. This suggests us to explore ways to combine the proposed 2D-ResNet18+LSTM and the 3D-ResNet18 methods to achieve high accuracies throughout the age range of interest. Our future work also includes more comprehensive evaluation in large multisite datasets and application to probing brain disorders in early life.

5. REFERENCES

- [1]. Cole James H, Ritchie Stuart J, Bastin Mark E, Valdés Hernández MC, Muñoz Maniega S, Royle Natalie, Corley Janie, Pattie Alison, Harris Sarah E, Zhang Qian, et al., “Brain age predicts mortality,” *Molecular psychiatry*, 2017.
- [2]. Cole James H, Poudel Rudra PK, Tsagkrasoulis Dimosthenis, Caan Matthan WA, Steves Claire, Spector Tim D, and Montana Giovanni, “Predicting brain age with deep learning from raw

- imaging data results in a reliable and heritable biomarker,” *NeuroImage*, vol. 163, pp. 115–124, 2017. [PubMed: 28765056]
- [3]. Aycheh Habtamu Minassie, Seong Joon-Kyung, Shin Jeong-Hyeon, Na Duk L, Kang Byungkon, Seo Sang Won, and Sohn Kyung-Ah, “Biological brain age prediction using cortical thickness data: a large scale cohort study,” *Frontiers in aging neuroscience*, vol. 10, pp. 252, 2018. [PubMed: 30186151]
- [4]. Ueda Masaru, Ito Koichi, Wu Kai, Sato Kazunori, Taki Yasuyuki, Fukuda Hiroshi, and Aoki Takafumi, “An age estimation method using 3D-CNN from brain MRI images,” in *ISBI*, 2019, pp. 380–383.
- [5]. Hochreiter Sepp and Schmidhuber Jürgen, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997. [PubMed: 9377276]
- [6]. Zhu Qikui, Du Bo, Turkbey Baris, Choyke Peter, and Yan Pingkun, “Exploiting interslice correlation for mri prostate image segmentation, from recursive neural networks aspect,” *Complexity*, Article 4185279, 2018.
- [7]. Evans Alan C, Brain Development Cooperative Group, et al., “The NIH MRI study of normal brain development,” *Neuroimage*, vol. 30, no. 1, pp. 184–202, 2006. [PubMed: 16376577]
- [8]. He Kaiming, Zhang Xiangyu, Ren Shaoqing, and Sun Jian, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [9]. Ulyanov Dmitry, Vedaldi Andrea, and Lempitsky Victor, “Instance normalization: The missing ingredient for fast stylization,” *arXiv preprint arXiv:1607.08022*, 2016.
- [10]. Rohlfing Torsten, Zahr Natalie M, Sullivan Edith V, and Pfefferbaum Adolf, “The sri24 multichannel atlas of normal adult human brain structure,” *Human brain mapping*, vol. 31, no. 5, pp. 798–819, 2010. [PubMed: 20017133]

Fig. 1:

The proposed framework for age estimation. Using NIH-PD as the Discovery Cohort, we train two different models (the typical 3D-CNN and the proposed 2D-CNN+LSTM) on the training set (80% of NIH-PD data) and test on the testing set (20% of NIH-PD data). Then, using MGHBCB as the Replication Cohort, we quantify the generalization ability across datasets.

Fig. 2:

Illustration of the 2D-ResNet with LSTM framework for age prediction. The 3D MRI is considered as a sequence of 2D images. The 2D-Resnet is used to extract features and then followed by a pooling operation. LSTM is used to capture the contextual information over the sequence and the final regressor is applied to estimate the age.

Fig. 3:

The structure of LSTM cell, which consists of a cell module c_t , a hidden state h_t , a forget gate f_t , an input gate i_t , and an output gate o_t .

Fig. 4: Cumulative score (CS) on the error levels from 0 to 2 years of different methods on the Discovery Cohort (NIH-PD testing set). A higher CS is better.

Fig. 5: Cumulative score (CS) on the error levels from 0 to 2 years trained on the Discovery Cohort (NIH-PD training set) and applied to the Replication Cohort. A high CS is better.

Table 1:

Accuracy in the Discovery Cohort. The MAE (in years) for brain age estimation on the NIH-PD testing set. Numbers in bold are more accurate results.

Method	Age group				Average
	0–5	6–10	11–15	16–22	
3D-ResNet18	0.30	1.12	1.19	1.85	1.07
2D-ResNet18+LSTM	0.32	0.92	1.02	1.79	0.96

Table 2:

Accuracy in the Replication Cohort, training on Discovery Cohort 0–22 years of age. The MAE (in years) of different methods for brain age estimation trained on the NIH-PD training set and replicated on the truly “unseen” MGHBCB dataset for examining the generalization ability. Numbers in bold are more accurate results.

Method	Age group						Average
	0–1	1–2	2–3	3–4	4–5	5–6	
3D-ResNet18	2.57	2.43	2.80	2.64	2.85	2.75	2.64
2D-ResNet18+LSTM	0.83	1.65	1.40	1.36	1.21	0.83	1.14

Table 3:

Accuracy in the Replication Cohort, training on Discovery Cohort 0–6 years of age. The MAE (in years) of different methods for brain age estimation trained on the NIH-PD training set with subjects of age 0–6 years and replicated on the truly “unseen” MGHBCB dataset for examining the generalization. Numbers in bold are more accurate results.

Method	Age group						Average
	0–1	1–2	2–3	3–4	4–5	5–6	
3D-ResNet18	0.66	0.48	0.43	0.75	1.26	2.10	0.79
2D-ResNet18+LSTM	0.24	1.34	1.43	1.25	0.69	0.92	0.78