



HHS Public Access

Author manuscript

Med Image Comput Comput Assist Interv. Author manuscript; available in PMC 2020 July 23.

Published in final edited form as:

Med Image Comput Comput Assist Interv. 2019 October ; 11765: 823–831.

doi:10.1007/978-3-030-32245-8_91.

Variational AutoEncoder For Regression: Application to Brain Aging Analysis

Qingyu Zhao¹, Ehsan Adeli¹, Nicolas Honnorat², Tuo Leng¹, Kilian M. Pohl^{1,2}

¹Stanford University, Stanford, CA, USA

²SRI International, Menlo Park, CA, USA

Abstract

While unsupervised variational autoencoders (VAE) have become a powerful tool in neuroimage analysis, their application to supervised learning is under-explored. We aim to close this gap by proposing a unified probabilistic model for learning the latent space of imaging data and performing supervised regression. Based on recent advances in learning disentangled representations, the novel generative process explicitly models the conditional distribution of latent representations with respect to the regression target variable. Performing a variational inference procedure on this model leads to joint regularization between the VAE and a neural-network regressor. In predicting the age of 245 subjects from their structural Magnetic Resonance (MR) images, our model is more accurate than state-of-the-art methods when applied to either region-of-interest (ROI) measurements or raw 3D volume images. More importantly, unlike simple feed-forward neural-networks, disentanglement of age in latent representations allows for intuitive interpretation of the structural developmental patterns of the human brain.

1 Introduction

Generative models in combination with neural networks, such as *variational autoencoders* (VAE), are often used to learn complex distributions underlying imaging data [1]. VAE assumes each training sample is generated from a latent representation, which is sampled from a prior Gaussian distribution through a neural-network, i.e., a decoder. Inferring the network parameters involves a variational procedure leading to an encoder network, which aims to find the posterior distribution of each training sample in the latent space. As an unsupervised learning framework, VAE has successfully been applied to several problems in neuroimaging, such as denoising [1], abnormality detection [2] or clustering tasks [3]. However, the use of VAE is still under-explored in the context of supervised regression; i.e., regression aims to predict a scalar outcome from an image based on a given set of training pairs. For instance in neuroimage analysis, the scalar could be a binary variable indicating if a subject belongs to the control or a disease group or a continuous variable encoding the age of a subject.

Several attempts have been made to integrate regression models into the VAE framework by directly performing regression analysis on the latent representations learned by the encoder

[4,5]. These works, however, still segregate the regression model from the autoencoder in a way that the regression needs to be trained by a separate objective function. To close the gap between the two models, we leverage recent advances in learning *disentangled latent representations* [6,7]. In the latent space, a representation is considered disentangled if changes along one dimension of that space are explained by a specific factor of variation (e.g., age), while being relatively invariant to other factors (e.g., sex, race) [6]. Herein, we adopt a similar notion to define a unified model combining regression and autoencoding. We then test the model for predicting the age of a subject solely based on its structural MR image.

Unlike a traditional VAE relying on a single latent Gaussian to capture all the variance in brain appearance, our novel generative age-predictor explicitly formulates the conditional distribution of latent representations on age while being agnostic to the other variables. Inference of model parameters leads to a combination between a traditional VAE network that models latent representations of brain images, and a regressor network that aims to predict age. Unlike the traditional VAE, our model is able to disentangle a specific dimension from the latent space such that traversing along that dimension leads to age-specific distribution of latent representations. We show that through this mechanism the VAE and the regressor networks regularize each other during the training process to achieve more accurate age prediction.

Next, we introduce the proposed VAE-based regression model in Section 2. Section 3 describes the experiments of age prediction for 245 healthy subjects based on their structural T1-weighted MR images. We implement the model using two network architectures: a multi-layer perception for imaging measurements and a convolutional neural network for 3D volume images. Both implementations achieve more accurate predictions compared to several traditional methods. Finally, we show that the learned age-disentangled generative model provides an intuitive interpretation and visualization of the developmental pattern in brain appearance, which is an essential yet challenging task in most existing deep learning frameworks.

2 VAE for Regression

Fig. 1 provides an overview of the model with blue blocks representing the generative model and red blocks the inference model.

The Generative Model

Let $\mathbf{X} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}\}$ be a training dataset containing structural 3D MR images of n subjects, and $\mathbf{C} = \{c^{(1)}, \dots, c^{(n)}\}$ be their age. We assume each MR image \mathbf{x} is associated with a latent representation $\mathbf{z} \in \mathbb{R}^M$, which is dependent on c . Then the likelihood distribution underlying each training image \mathbf{x} is $p(\mathbf{x}) = \int_{\mathbf{z}, c} p(\mathbf{x}, \mathbf{z}, c)$, and the generative process of \mathbf{x} reads $p(\mathbf{x}, \mathbf{z}, c) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z}|c)p(c)$, where $p(c)$ is a prior on age. In a standard VAE setting [8], the ‘decoder’ $p(\mathbf{x}|\mathbf{z})$ is parameterized by a neural network f with the generative parameters θ , i.e., $p(\mathbf{x}|\mathbf{z}) \sim \mathcal{N}(\mathbf{x}; f(\mathbf{z}; \theta), \mathbf{I})$ ³. Different from the traditional VAE is the modeling of latent representations. Instead of using a single Gaussian prior to generate \mathbf{z} ,

we explicitly condition \mathbf{z} on age c , such that the conditional distribution $p(\mathbf{z}|c)$ captures an age-specific prior on latent representations. We call $p(\mathbf{z}|c)$ a *latent generator*, from which one can sample latent representations for a given age. We further assume the nonlinearity of this generative model can be fully captured by the decoder network $p(\mathbf{x}|\mathbf{z})$, such that a linear model would suffice to parameterize the generator: $p(\mathbf{z}|c) \sim \mathcal{N}(\mathbf{z}; \mathbf{u}^T c, \sigma^2 \mathbf{I})$, $\mathbf{u}^T \mathbf{u} = 1$. With this construction we can see that \mathbf{u} is essentially the disentangled dimension [6] associated with age; traversing along \mathbf{u} yields age-specific latent representations. Note this model does not reduce the latent space to 1D but rather links one dimension of the space to age.

Inference Procedure

The parameters of the above generative model can be estimated via maximum likelihood estimation (MLE) i.e., by maximizing the sum of log likelihood $\sum_{i=1}^N \log p(\mathbf{x}^{(i)})$. For such an optimization, we adopt a standard procedure of variational inference and introduce an auxiliary function $q(\mathbf{z}^{(i)}, c^{(i)}|\mathbf{x}^{(i)})$ to approximate the true posterior $p(\mathbf{z}^{(i)}, c^{(i)}|\mathbf{x}^{(i)})$. In the following we omit index i for convenience. In so doing, $\log p(\mathbf{x})$ can be rewritten as the sum of the KL-divergence D_{KL} between $q(\mathbf{z}, c|\mathbf{x})$ and $p(\mathbf{z}, c|\mathbf{x})$ and the ‘variational lower-bound’ $\mathcal{L}(\mathbf{x})$:

$$\log p(\mathbf{x}) = D_{KL}(q(\mathbf{z}, c|\mathbf{x}) \parallel p(\mathbf{z}, c|\mathbf{x})) + \mathcal{L}(\mathbf{x}). \quad (1)$$

Based on the mean-field theory, we further assume $q(\mathbf{z}, c|\mathbf{x}) = q(\mathbf{z}|\mathbf{x})q(c|\mathbf{x})$.

Then the lower-bound can be derived as

$$\begin{aligned} \mathcal{L}(\mathbf{x}) := & -D_{KL}(q(c|\mathbf{x}) \parallel p(c)) \\ & + \mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p(\mathbf{x}|\mathbf{z})] - \mathbb{E}_{q(c|\mathbf{x})}[D_{KL}(q(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z}|c))] \end{aligned} \quad (2)$$

In the above equation, we formulate $q(c|\mathbf{x})$ as a univariate Gaussian $q(c|\mathbf{x}) \sim \mathcal{N}(c; f(\mathbf{x}; \phi_c), g(\mathbf{x}; \phi_c)^2)$, where ϕ_c are the parameters of the inference networks. We can see that $q(c|\mathbf{x})$ is essentially a regular feed-forward regression network with an additional output being the uncertainty (i.e., standard deviation) of the prediction. In this work we call $q(c|\mathbf{x})$ a *probabilistic regressor*. In an unsupervised setting, the KL-divergence of Eq. (2) regularizes the prediction of c with a prior. However, in our supervised setting this term can be simply replaced by $\log q(c|\mathbf{x})$ as the ground-truth of c is known for each training sample [9,10]⁴.

Similar to a traditional VAE, the remaining part of the inference involves the construction of a *probabilistic encoder* $q(\mathbf{z}|\mathbf{x})$, which maps the input image \mathbf{x} to a posterior multivariate Gaussian distribution in the latent space $q(\mathbf{z}|\mathbf{x}) \sim \mathcal{N}(\mathbf{z}; f(\mathbf{x}; \phi_z), g(\mathbf{x}; \phi_z)^2 \mathbf{I})$. Then the second term of Eq. (2) encourages the decoded reconstruction from the latent representation to resemble the input [8]. The third term of of Eq. (2) encourages the posterior $q(\mathbf{z}|\mathbf{x})$ to

³when \mathbf{x} is binary, a Bernoulli distribution can define $p(\mathbf{x}|\mathbf{z}) \sim \text{Ber}(\mathbf{x}; f(\mathbf{z}; \theta))$

⁴In a semi-supervised setting where no informative prior in present, $\mathbb{H}(q(c|\mathbf{x}))$, i.e., the entropy of $q(c|\mathbf{x})$, is commonly used to replace the last term of Eq. 4 for samples with unknown c [9,10].

resemble the age-specific prior $p(z|c)$. This is the main mechanism for linking latent representations with age prediction: on the one hand, latent representations generated from the predicted c have to resemble the latent representation of the input image and on the other hand, age-linked variation in the latent space is encouraged to follow a direction defined by u . We used the SGVB estimator with the reparametrization trick [8] to optimize the expectation in the last two terms of Eq. (2).

Lastly, it has been shown that the supervised training of end-to-end feed-forward neural networks often suffers from over-fitting problems, whereas unsupervised autoencoders can often learn robust and meaningful intermediate features that are transferable to supervised tasks [11]. By combining the two frameworks, our model therefore allows for the sharing of low-level features (e.g., by convolutional layers) jointly learned by the autoencoder and regressor.

3 Experiments

Understanding structural changes of the human brain as part of normal aging is an important topic in neuroscience. One emerging approach for such analysis is to learn a model that predicts age from brain MR images and then to interpret the patterns learned by the model. We tested the accuracy of the proposed regression model in predicting age from MRI based on two implementations⁵: the first implementation was based on a multi-layer perceptron neural network (all densely connected layers) applied to ROI-wise brain measurements while the second implementation was based on convolutional neural networks (CNN) applied to 3D volume images focusing on the ventricular area. Both implementations were cross-validated on a dataset consisting of T1-weighted MR images of 245 healthy subjects (122/123 women/men; ages 18 to 86) [12]. There was no group-level age difference between females and males ($p = 0.51$, two-sample t -test).

3.1 Age Prediction Based on ROI Measures

With respect to the perceptron neural network, the input of the encoder were the z-scores of 299 ROI measurements generated by applying FreeSurfer (V 5.3.0) to the skull-stripped MR image of each subject [12]. The measurements consisted of the mean curvature, surface area, gray matter volume, and average thickness of 34 bilateral cortical ROI, the volumes of 8 bilateral subcortical ROIs, the volumes of 5 subregions of the corpus callosum, the volume of all white matter hypointensities, the left and right lateral and third ventricles, and the supratentorial volume (svol).

The input to the encoder was first densely connected to 2 intermediate layers of dimension (128,32) with tanh as the activation function. The resulting feature was then separately connected to two layers of dimension 8 yielding the mean and diagonal covariance of the latent representation. The regressor shared the 2 intermediate layers of the encoder (Fig. 1b) and produced the mean and standard deviation for the predicted age. The decoder took a latent representation as input and adopted an inverse structure of the encoder for reconstruction.

⁵Implementation based on Tensorflow 1.7.0, keras 2.2.2. Source code available at <https://github.com/QingyuZhao/VAE-for-Regression>

3.2 Age Prediction Based on 3D MR Images

Taking advantage of recent advances in deep learning, the second implementation was built upon convolutional neural networks that directly took 3D images as input. All skull-stripped T1 images were registered to the SRI24 atlas space and down-sampled to 2mm isotropic voxel size. Since it has been well established that the ventricular volume significantly increases with age [13], we then tested if the model would learn this structural change in predicting age. To do this, each image was cropped to a 64*48*32 volume containing the ventricle region and was normalized to have zero mean and unit variance. This smaller field of view allowed for faster and more robust training of the following CNN model on limited sample size (N=245). Specifically, the encoder consisted of 3 batches of 3*3*3 convolutional layers with rectified linear unit (ReLU) activation and 2*2*2 max pooling layers. The sizes of feature banks in the convolutional layers were (16,32,64) respectively.

Similar to the previous implementation, the extracted features were fed into 2 densely connected layers of dimension (64,32) with tanh activation function. The final dimension of latent space was 16. The regressor shared the convolutional layers of the encoder and also had 2 densely connected layers of (64,32). The decoder had an inverse structure of the encoder and used Upsampling3D as the inverse operation of max pooling. Since the CNN-based implementation had substantially more model parameters to determine than the first implementation, L2 regularization was applied to all densely connected layers.

3.3 Measuring Accuracy

The accuracy of each implementation was reported based on a 5-fold cross-validation measuring the R2 score (coefficient of determination, the proportion of the variance in age that is predictable from the model) and root mean squared error (rMSE). The outcome of each approach was compared to 7 other regression methods, of which 6 were non-neural-network methods as implemented in *scikit-learn* 0.19.1: linear regression (LR), Lasso, Ridge regression (RR), support vector regression (SVR), gradient-boosted tree (GBT), k-nearest neighbour regression (K-NN). The last approach was a single neural-network regressor (NN), i.e., the component corresponding to $q(\mathbf{c}; \mathbf{x})$ without outputting standard deviation.

With respect to the ROI-based experiments, optimal hyperparameters of the scikit-learn methods (except for LR) were determined through a 10-fold inner cross-validation (an overall nested cross-validation). Specifically, we searched $C \in \{1, 10, \dots, 10^3\}$, $\gamma \in \{10^{-2}, \dots, 10^2\}$ for SVR, $N \in \{10, 50, 100, 500\}$ for GBT, $\alpha \in \{10^{-3}, \dots, 10^4\}$, $\gamma \in \{10^{-2}, \dots, 10^2\}$ for RR, and $N \in \{1, 5, 10, 50\}$ for K-NN.

With respect to the 3D-image-based experiments, nested cross-validation was extremely slow for certain methods (e.g. GBT, CNN), so we simply repeated the outer 5-fold cross-validation using the hyperparameters defined in the above search space and reported the best accuracy. The search space of the L2 regularization for NN and our method was $\{0, .001, .01, .1, 1\}$.

3.4 Results

As Table 1 shows, age prediction based on 3D images of ventricle was generally more accurate than on ROI measurements. The two neural-network-based predictions were the most accurate in terms of R2 and rMSE. Fig. 3 shows the predicted age (in the 5 testing folds) estimated by our model versus ground-truth. The best prediction was achieved by our model applied to the 3D ventricle images, which yielded a 6.9-year rMSE. In the ROI-based experiment, our model was more accurate than the single neural-network regressor (NN), which indicates the integration of VAE for modeling latent representations could regularize the feed-forward regressor network. In the 3D-image-based experiment, our model was more accurate than NN either with (Table 1) or without L2 regularization (our model and NN yielded R2 of 0.761 and 0.745 respectively). Even though this improvement was not as significant as in the ROI-based experiment, our model enabled direct visualization of brain developmental patterns.

Indeed, despite the tremendous success of deep learning in various applications, interpretability of the black-box CNN (e.g., which input variable leads to accurate prediction, or what specific features are learned) remains an open research topic. Most existing solutions can only produce a “heat map” indicating the location of voxels that contribute to faithful prediction, but this does not yield any semantic meaning of the learned features that can improve mechanistic understanding of the brain. Thanks to the generative modelling, our formulation provides an alternative way for interpreting the aging pattern captured by the CNN. Specifically, Fig. 2 shows the simulated “mean brain images” at different ages by decoding age-specific latent representations $\{z = \mathbf{u}^T c | c \in [18, 86]\}$, i.e., mean of the latent generator $p(z|c)$. We can clearly observe that the pattern learned by the model for age prediction was mainly linked to the enlargement of ventricle. This result is consistent with current understanding of the structural development of the brain.

Lastly, we show in Figure 3 that the dimension related to age was disentangled from the latent space. In both ROI-based and image-based experiments, we trained our model on the entire dataset. The resulting latent representations were transformed from the latent space to a 2D plane via TSNE and color-coded by the ground-truth age. We observe that one direction of variation is associated with age, whereas the unsupervised training of traditional VAE does not lead to clear disentanglement.

4 Conclusion and Discussion

In this paper, we introduced a generic regression model based on the variational autoencoder framework and applied it to the problem of age prediction from structural MR images. The novel generative process enabled the disentanglement of age as a factor of variation in the latent space. This did not only produce more accurate prediction than a regular feed-forward regressor network, but also allowed for synthesizing age-dependent brains that facilitated the identification of brain aging pattern. Future direction of this work includes simultaneously disentangling more demographics factors of interest, e.g. sex, disease group, to study compounding effects, e.g. age-by-sex effects or accelerated aging caused by disease.

Acknowledgements

This research was supported in part by NIH grants U24AA021697, AA005965, AA013521, AA017168, AA026762

References

1. Benou A, Veksler R, Friedman A, Raviv TR: De-noising of contrast-enhanced mri sequences by an ensemble of expert deep neural networks. In: International Workshop on Deep Learning in Medical Image Analysis (2016)
2. Baur C, Wiestler B, Navab N: Deep autoencoding models for unsupervised anomaly segmentation in brain mr images. In: MICCAI (2018)
3. Zhao Q, Honnorat N, Adeli E, Pfefferbaum A, Sullivan EV, Pohl KM: Variational autoencoder with truncated mixture of gaussians for functional connectivity analysis. In: Information Processing in Medical Imaging (2019) 867–879
4. Yoo Y, et al.: Variational autoencoded regression: High dimensional regression of visual data on complex manifold. In: CVPR (2017)
5. Chen L, et al.: Deep autoencoding models for unsupervised anomaly segmentation in brain mr images. Information Sciences 428 (2018) 49–61
6. Higgins I, et al.: Beta-vae: Learning basic visual concepts with a constrained variational framework. In: ICLR (2017)
7. Kim H, Mnih A: Disentangling by factorising. In: ICML (2018)
8. Kingma D, Welling M: Auto-encoding variational bayes. In: ICLR (2013)
9. Kingma DP, Rezende DJ, Mohamed S, Welling M: Semi-supervised learning with deep generative models. In: NeurIPS (2014)
10. Nalisnick E, Smyth P: Stick-breaking variational autoencoders. In: ICLR (2017)
11. Zhuang F, Cheng X, Luo P, Pan SJ, He Q: Supervised representation learning: Transfer learning with deep autoencoders. In: IJCAI (2015)
12. Adeli E, Kwon D, Zhao Q, Pfefferbaum A, Zahr N, Sullivan E, Pohl K: Chained regularization for identifying brain patterns specific to HIV infection. Neuroimage 183 (2018) 425–437 [PubMed: 30138676]
13. Kaye J, DeCarli C, Luxenberg J, Rapoport S: The significance of age-related enlargement of the cerebral ventricles in healthy men and women measured by quantitative computed x-ray tomography. J Am Geriatr Soc. 40(3) (1992) 225–31 [PubMed: 1538040]

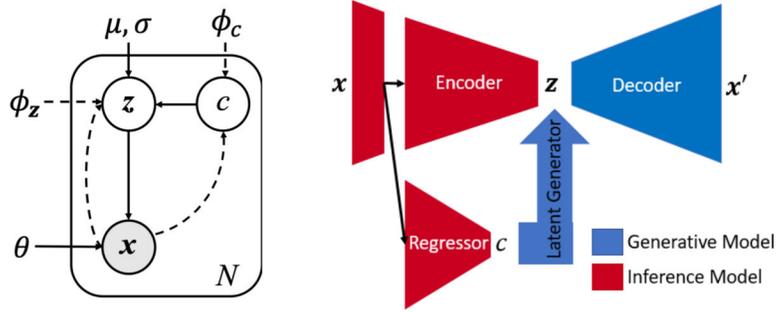


Fig.1. Probabilistic (left) and graphical (right) diagrams of the VAE-based regression model. Each image x is assumed to be generated from its representation z , which is dependent on age c (blue blocks). The inference model (red blocks) constructs a probabilistic encoder for determining the latent representation and a probabilistic regressor for predicting age.

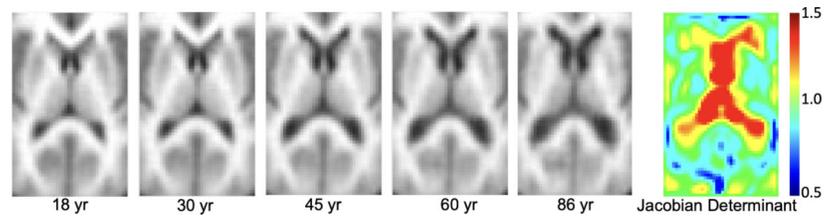


Fig.2.

Left: Brain images reconstructed from age-specific latent representations. Right: Jacobian determinant map derived from the registration between the 18 year old brain and the 86 year old brain. The major expanding region is located on the ventricle.

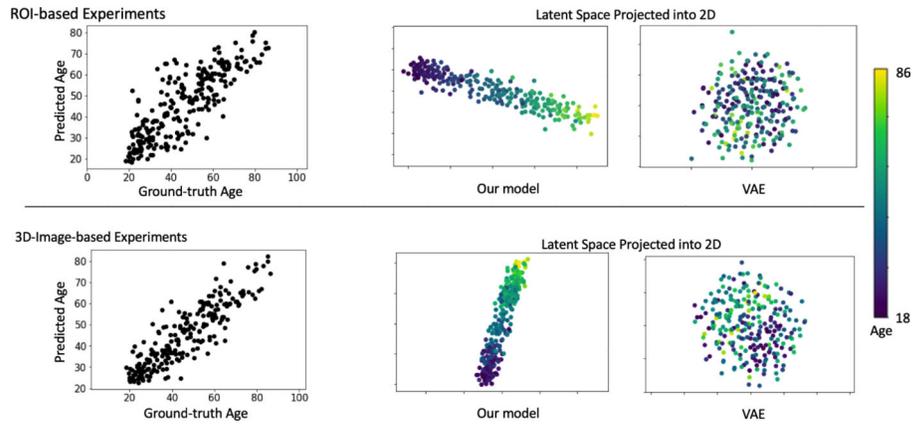


Fig.3. Upper row: results of ROI-based experiments. Lower row: results of 3D-imagebased experiments. Left: Predictions made by our model vs. ground-truth. Middle: Latent representations estimated by our model. Right: Latent representations estimated by traditional VAE.

Table 1.

Age prediction accuracy of different methods based on ROI measurements and 3D volume images of ventricle.

		LR	RR	SVR	GBT	K-NN	NN	Ours
ROI Measures	R2	0.107	0.336	0.311	0.64	0.535	0.563	0.666
	rMSE	14.6	12.6	12.8	9.3	10.5	10.3	9.0
3D Volume	R2	0.737	0.737	0.737	0.719	0.549	0.79	0.808
	rMSE	7.8	7.8	7.8	8.2	10.5	7.0	6.9

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript