



Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Longitudinal self-supervised learning

Qingyu Zhao^{a,1}, Zixuan Liu^{b,1}, Ehsan Adeli^{a,c}, Kilian M. Pohl^{a,d,*}^a Department of Psychiatry & Behavioral Sciences, Stanford University, Stanford, CA 94305, USA^b Department of Electrical Engineering, Stanford University, Stanford, CA 94305, USA^c Department of Computer Science, Stanford University, Stanford, CA 94305, USA^d Center for Biomedical Sciences, SRI International, Menlo Park, CA 95025, USA

ARTICLE INFO

Article history:

Received 29 November 2020

Revised 19 February 2021

Accepted 24 March 2021

Available online 4 April 2021

Keywords:

Longitudinal neuroimaging

Self-supervised learning

Factor disentanglement

Brain age

ABSTRACT

Machine learning analysis of longitudinal neuroimaging data is typically based on supervised learning, which requires large number of ground-truth labels to be informative. As ground-truth labels are often missing or expensive to obtain in neuroscience, we avoid them in our analysis by combining factor disentanglement with self-supervised learning to identify changes and consistencies across the multiple MRIs acquired of each individual over time. Specifically, we propose a new definition of disentanglement by formulating a multivariate mapping between factors (e.g., *brain age*) associated with an MRI and a latent image representation. Then, factors that evolve across acquisitions of longitudinal sequences are disentangled from that mapping by self-supervised learning in such a way that changes in a single factor induce change along one direction in the representation space. We implement this model, named Longitudinal Self-Supervised Learning (LSSL), via a standard autoencoding structure with a cosine loss to disentangle brain age from the image representation. We apply LSSL to two longitudinal neuroimaging studies to highlight its strength in extracting the brain-age information from MRI and revealing informative characteristics associated with neurodegenerative and neuropsychological disorders. Moreover, the representations learned by LSSL facilitate supervised classification by recording faster convergence and higher (or similar) prediction accuracy compared to several other representation learning techniques.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

A longitudinal study repeatedly measures the same variable to track a specific group of individuals over a period of time (Caruana et al., 2015). For example, longitudinal neuroimaging studies are often used to evaluate the impact of age on the brain (Dennis and Cabeza, 2008), the relationship between risk factors and development of disease (Jack et al., 2008), and the outcomes of treatments over time (Liu et al., 2010). A critical component of longitudinal studies is to apply data analysis approaches properly modeling the complex correlations underlying the repeated measures, which are often characterized by a mixture of inter-subject variance and intra-subject dependency. Popular analysis approaches are mixed-effect models (Bernal-Rusiel et al., 2013) and analysis of variance (ANOVA) (Cash et al., 2016), which can inspect the influence of key factors (e.g., age or disease) on individual imaging measurements (e.g., cortical thickness of regions of in-

terest) extracted from the longitudinal Magnetic Resonance images (MRIs). However, this type of univariate analysis ignores the high-dimensional relations across multiple brain regions (Habeck and Stern, 2010). With recent advances in deep learning, this limitation can be largely resolved by data-driven supervised learning, i.e., by training models to predict the value of a set of factors (e.g., age or diagnosis group) for each subject directly from their raw images (Aghili et al., 2018).

One limitation of supervised learning is that the training requires large amount of data with accurate labels, which is infeasible for some neuroimaging applications. For example, prior studies have trained models to predict subjects' age from their structural MRIs to understand the effect of aging on brain morphometry (Smith et al., 2020; Zhao et al., 2019). However, such age-based supervision can be sub-optimal as the chronological age does not always reflect the true condition of aging in the brain (Steffener et al., 2016). To resolve this issue, we propose here a novel learning scheme that replaces the direct supervision with *self-supervision* in the context of repeated measures. In general, self-supervised learning (Kolesnikov et al., 2019) aims to automatically generate *supervisory signal* by exploring similarity and dissimilarity relations across samples through the learning process of

* Corresponding author at: Department of Psychiatry & Behavioral Sciences, Stanford University, Stanford, CA 94305, USA.

E-mail address: kilian.pohl@stanford.edu (K.M. Pohl).

¹ Authors contributed equally.

their representations. This concept aligns with the longitudinal design, in which each subject serves as their own ‘comparison’ with respect to the change over time. Therefore, the repeated measures could fully leverage the structured inter-relation across time points to learn time-dependent representations of MRIs.

Another concept in representation learning that can be related to the longitudinal design is *factor disentanglement* (Tschannen et al., 2018). The key intuition behind disentanglement is that real-world data have low-dimensional representations that are controlled by distinct and interpretable factors (Higgins et al., 2018). Changing a single factor should thus lead to a change in a single dimension in the representation space. According to a recent publication (Locatello et al., 2018), factor disentanglement cannot be stringently formulated for cross-sectional data without explicitly supervision. However, in a longitudinal study, where each individual is repeatedly examined over time, all time-independent factors are fixed so that the impact of time-related factors, such as age, can be effectively revealed on the observations. In other words, the study design of repeated measures fits perfectly to ‘supervising’ the disentanglement of the evolving factor across the measures.

Inspired by these affinities, we propose a representation learning strategy, named Longitudinal Self-Supervised Learning (LSSL), to investigate the impact of aging and neuropsychological disorders on the brain within the context of longitudinal MRIs. Thanks to the repeated measure design in longitudinal studies, we separately define the space of factors and the representation space of MRIs and formulate disentanglement in the context of multivariate mapping between the two spaces. The factor of brain age is then disentangled by self-supervision on the temporal order between the multiple MRIs for each subject. The training uses this ordinal information to encourage the change of subject-specific representations to follow a common developmental direction. This objective is condensed to optimizing a combination of image reconstruction loss of a standard autoencoder (Chen et al., 2017) and a simple cosine loss in the representation space.

We test our model on a synthetic dataset and two longitudinal MR data sets: one investigates the impact of Alzheimer’s Disease (AD) on the brain and the other one of alcohol dependence. While the training does not rely on any ground-truth age nor diagnostic labels, LSSL successfully disentangles a factor in the representation space linked to brain age, which is superior than the chronological age in characterizing the health condition of the brain. In addition, LSSL successfully reveals accelerated aging effects of AD and alcohol dependence compared to the control cohort. When performing a downstream task of predicting diagnosis labels of subjects, the representations and pre-trained encoder learned by our model result in faster convergence and more (or equivalently) accurate classification accuracy compared to several commonly used state-of-the-art unsupervised or self-supervised representation learning strategies.

2. Related work

2.1. Longitudinal neuroimaging studies

Deep learning models applied to longitudinal neuroimaging studies are largely based on supervised learning, i.e., by training classification models based on image time series using regular recurrent networks (Lipton et al., 2015; Santeramo et al., 2018; Cui et al., 2019; Ghazi et al., 2019). Based on these models, longitudinal pooling is proposed to augment the learned representation of each time point with information gathered from images of other time points in the series (Ouyang et al., 2021). Other methods for explicitly exploiting dependencies within the intra-subject series are generally based on parameterizing the trajectories of representa-

tions in the latent space, such as using Mixed Effect Models (Xiong et al., 2019; Louis et al., 2019).

2.2. Brain-age analysis

The biological ‘brain age’ is a marker quantifying the health condition of the brain and correlates with mental and physical fitness (Steffener et al., 2016). To predict the brain age from structural MRIs, recent publications (Cole et al., 2016; Elliott et al., 2019; Kaufmann et al., 2019; Smith et al., 2020; Zhao et al., 2019) proposed to train supervised learning models on a cohort of healthy subjects with the prediction target being their chronological age. Such models can be built on imaging measurements (Smith et al., 2020) or raw 3D images (Zhao et al., 2019) and in both cross-sectional (Smith et al., 2020) and longitudinal settings (Elliott et al., 2019). After training, the model can be applied to both healthy and diseased subjects, and the difference between the predicted brain age and chronological age can be regarded as a phenotype related to the impact of the brain disorder on normal aging (Kaufmann et al., 2019). Despite the meaningful findings revealed in these studies, one common pitfall is the underlying assumption that the chronological age exactly corresponds to the brain age, which is challenged by several recent studies that show brain aging is highly heterogeneous even within the healthy populations (Franke and Gaser, 2019).

2.3. Self-supervised learning

To tackle the problem of missing or expensive-to-obtain ground-truth labels required by supervised learning, self-supervised methods learn representations by training models in tasks that do not require explicit annotations or labels (Jing and Tian, 2019), such as image inpainting (Zhang et al., 2016), colorization (Larsson et al., 2017), super-resolution (Dong et al., 2014; Johnson et al., 2016), and predicting spatial relationships between two image patches from the same image (Noroozi and Favaro, 2016; Sabokrou et al., 2019). Beyond these within-sample learning schemes, contrastive learning is a self-supervised approach that explicitly models the relationships across samples, such as by distinguishing between similar and dissimilar images (Chen et al., 2020), modeling temporal dependency across time (van den Oord et al., 2018), and estimating depth from stereo images (Pillai et al., 2019). Once trained, the resulting representations can be embedded into supervised learning tasks, such as multi-task and cross-domain feature learning (Doersch and Zisserman, 2017; Ren and Lee, 2018), which result in a more efficient training with respect to labelled data and computational resources than training the supervised models from scratch (Chen et al., 2020; He et al., 2019; Kolesnikov et al., 2019).

2.4. Factor disentanglement

While there is no consensus on the mathematical definition of *disentanglement*, conceptually, a representation is considered disentangled if changes along one dimension of the representation are explained by a specific factor of variation (e.g., age) while being relatively invariant to other factors (e.g. gender, ethnicity) (Higgins et al., 2018). Most existing works formulate this notion from a statistical perspective by pursuing statistical independence among random variables in the latent space (factorizable latent representations Tschannen et al., 2018; Higgins et al., 2018). Therefore, state-of-the-art approaches for unsupervised disentanglement learning are based on a Variational Autoencoder (VAEs) structure, which aims to learn a factorizable posterior from the marginal distribution of the observed data (Higgins et al., 2017; Chen et al., 2018; Kim and Mnih, 2018; Zhao et al., 2017). Despite promising

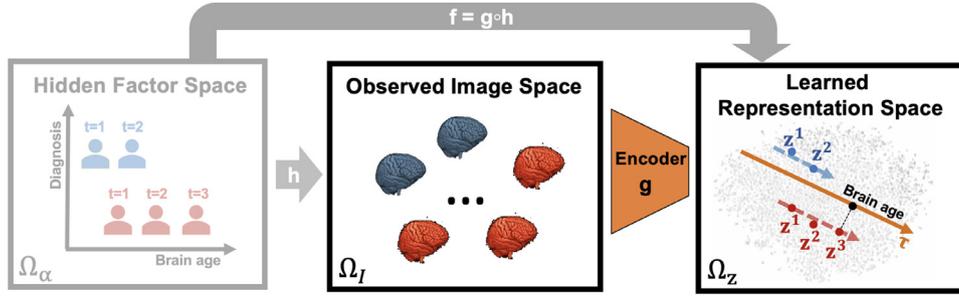


Fig. 1. Longitudinal Self-Supervised Learning (LSSL) aims to learn representations from observed images, which are assumed to be generated from a set of hidden factors. In this example, the variation of the repeated measures of two subjects (blue and red, t encodes the order of visits) is assumed to relate to an increase in brain age. LSSL then disentangles a 1D direction τ linked to brain age from the representation space such that the developmental trajectories of subject-specific representations \mathbf{z}^i are colinear with τ . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

results, the study by Locatello et al. (2018) challenges the theoretical validity of this idea. They point out that given any marginal distribution of the observed data, there exists an infinite number of generative processes from either disentangled or fully entangled latent representations. Therefore, a true factor disentanglement requires supervision, for which we propose to perform self-supervision on repeated measures.

3. Longitudinal self-supervised learning

We now first provide a new perspective of factor disentanglement by defining a multivariate mapping from a hidden factor space underlying the observed images to a representation space learned from those images (see Fig. 1). This setup motivates us to define a novel self-supervised objective function that does not depend on the statistical property of the disentangled factor. We then employ this formulation in the context of the repeated measures design to disentangle brain age from longitudinal MRI data.

3.1. Disentangled mapping from factor to representation space

Current unsupervised learning frameworks often define the concept of *disentanglement* as the observed data having factorizable distributions in the representation space, such that each representation variable corresponds to a real-world factor that is statistically independent from other variables (Higgins et al., 2017). This learning objective is impractical because factors in real world datasets are not necessarily statistically independent. For example, brain morphology is influenced by both gender and brain size, two highly correlated factors. Thus, deriving a 2D representation of the underlying distribution could yield two statistically independent directions (such as with Principal or Independent Component Analysis), but the meaning of each direction with respect to real-world factors would be unclear. To resolve this issue, we step away from the statistical formulation and approach the disentanglement problem from a pure algebraic perspective. Specifically, we propose to explicitly separate the concepts of the *space of factors* from the representation space. We do so by assuming images are generated by factors, and hence can be reduced to a low-dimensional representation (see Fig. 1). We then are interested in the deterministic multivariate mapping function between the two spaces as a means of disentanglement.

To mathematically formulate such mappings, we denote $\Omega_\alpha = \mathbb{R}^M$ as the *factor space*, and $\Omega_I = \mathbb{R}^P$ the *image space*. We assume each image $\mathbf{I} \in \Omega_I$ can be fully determined based on M factors $\alpha := [\alpha_1, \dots, \alpha_M] \in \Omega_\alpha$ through a differentiable generative process $\mathbf{I} = \mathbf{h}(\alpha)$. Further, we aim to learn a differentiable encoder $\mathbf{g}(\cdot)$ that reduces the image \mathbf{I} to a K dimensional representation $\mathbf{z} := [z_1, \dots, z_K]$ in the *representation space* Ω_Z , i.e., $\mathbf{z} =$

$\mathbf{g}(\mathbf{I}) \in \Omega_Z$ and $z_i = g_i(\mathbf{I})$. We then define $\mathbf{f} : \Omega_\alpha \rightarrow \Omega_Z$ as the composite multivariate-to-multivariate mapping $\mathbf{f} := \mathbf{g} \circ \mathbf{h}$, where $\mathbf{f} = [f_1, \dots, f_K]$ and each f_i is a differentiable multivariate-to-univariate mapping with $z_i := f_i(\alpha)$ for $i \in [1, K]$.

Without loss of generality, we assume that z_1 is linked to the first factor α_1 . Let $\hat{\alpha} := [\alpha_2, \dots, \alpha_M]$, we then consider the factor α_1 as disentangled within the representation \mathbf{z} if \mathbf{f} can be factorized such that

$$z_1 = f_1(\alpha_1) \text{ with } f_1 : \mathbb{R} \rightarrow \mathbb{R} \text{ being strictly increasing, and} \\ z_i = f_i(\hat{\alpha}) \text{ for } i > 1. \quad (1)$$

In other words, disentanglement is achieved when (1) z_1 is solely dependent on α_1 (the monotonicity of f_1 ensures the mapping to be bijective, i.e., without loss of information, and preserve ordinal information of factor z_1); and (2) the remaining representation in \mathbf{z} is solely dependent on factors other than α_1 .

3.2. Self-supervised disentanglement

In many applications, the only available data are a set of images \mathbf{I} . The underlying generative process, including the mapping function \mathbf{h} , dimension of factor space M , and values of α are hidden from observation, so training disentangled representation with respect to a specific factor can become extremely challenging. However, in situations where each training sample has multiple images measured with respect to different values of a specific factor, we can leverage self-supervision to achieve disentanglement. To show this, the factorization of \mathbf{f} in Eq. (1) can be transformed to the following conditions:

$$\frac{\partial f_1}{\partial \alpha_1} > 0, \text{ and } \frac{\partial f_i}{\partial \alpha_1} = 0, \frac{\partial f_1}{\partial \alpha_i} = 0 \text{ for } i > 1. \quad (2)$$

Based on the chain rule, the partial derivative of f_i with respect to α_j can be further transformed to a directional derivative with respect to the vector \mathbf{u}_j :

$$\mathbf{u}_j = \left[\frac{\partial h_1}{\partial \alpha_j}, \dots, \frac{\partial h_p}{\partial \alpha_j} \right], \quad (3)$$

$$\frac{\partial f_i}{\partial \alpha_j} = \sum_{p=1}^P \frac{\partial f_i}{\partial h_p} \frac{\partial h_p}{\partial \alpha_j} = \nabla_{\mathbf{u}_j} g_i(\mathbf{I}). \quad (4)$$

Eq. (4) translates the problem setup on \mathbf{f} and α to a setup with respect to \mathbf{g} and \mathbf{I} . Specifically, let \mathbf{u}_j be the change in the image space after perturbing the value of α_j by ϵ and $\mathbf{I}' := \mathbf{I} + \mathbf{u}_j$, then the corresponding change in z_i can be defined as

$$\frac{\partial f_i}{\partial \alpha_j} \approx \frac{g_i(\mathbf{I}') - g_i(\mathbf{I})}{\epsilon}. \quad (5)$$

As such, the disentanglement defined by Eq. (2) is achieved if and only if

Condition 1: Upon perturbation of α_1 , $\mathbf{g}(\mathbf{I}^t) - \mathbf{g}(\mathbf{I})$ is co-linear with $[1, 0, \dots, 0]$;

Condition 2: Upon perturbation of α_i for $i > 1$, $\mathbf{g}_i(\mathbf{I}^t) - \mathbf{g}_i(\mathbf{I}) = 0$.

In other words, to disentangle α_1 from the representation space, one should optimize for an encoder \mathbf{g} such that the two conditions apply to the representations of all pairs of images. This objective is self-supervised as we only need to pair the images but not to provide the true values of the factor.

3.3. Studying brain aging via longitudinal MRIs

As mentioned, brain age characterizes the apparent health condition of the brain but not necessarily equals chronological age. For example, a patient with neurodegenerative disease can have a higher brain age than a healthy subject albeit they have the same chronological age (Elliott et al., 2019; Kaufmann et al., 2019). This section describes how to leverage the prior self-supervised learning model to disentangle brain age from longitudinal structural MRIs.

To do so, we assume that brain age is the dominant factor that changes the brain morphology of an individual across the longitudinal scans while the other genotypic and demographic factors (such as gender and ethnicity) are static over time. Based on this assumption, *Condition 2* is omitted from the following analysis as one can not ‘perturb’ those static factors to examine their influence on the image representations in a longitudinal design. As such, we now present the Longitudinal Self-Supervised Learning (LSSL) algorithm that performs disentanglement guided by *Condition 1*.

Let \mathcal{I} be the collection of all MR images and \mathcal{S} be the set of subject-specific image pairs; i.e., \mathcal{S} contains all $(\mathbf{I}^t, \mathbf{I}^s)$ such that \mathbf{I}^t and \mathbf{I}^s are from the same subject with \mathbf{I}^t scanned before \mathbf{I}^s . To apply *Condition 1* to disentangle brain age from an image pair, we first assume that the small increase of brain age between the two time points corresponds to the perturbation of α_1 . Next, we relax the colinearity constraint of *Condition 1* that the disentangled direction has to align with the first natural coordinate axis in the representation space (i.e., $[1, 0, \dots, 0]$). Instead, we parameterize the direction linked to brain age as a free-form 1D unit vector $\boldsymbol{\tau} \in \Omega_z$ that can be jointly learned during training (Fig. 1). This strategy is motivated by the findings of Rolinek et al. (2018), which suggest that the encoder network by itself does not have the capacity to model arbitrary rotations of the representation space.

In doing so, the brain age associated with the two images, ψ^t and ψ^s , can be defined as the projections of the corresponding image representations to $\boldsymbol{\tau}$, e.g., $\psi^t = \mathbf{g}(\mathbf{I}^t)^\top \boldsymbol{\tau} = \mathbf{z}^t \boldsymbol{\tau}$ (Fig. 1). Then, ensuring *Condition 1* while preserving $\psi^s > \psi^t$ is equivalent to enforcing $\cos(\mathbf{g}(\mathbf{I}^s) - \mathbf{g}(\mathbf{I}^t), \boldsymbol{\tau}) = 1$, i.e., a zero-angle between $\boldsymbol{\tau}$ and the direction of progression in the representation space. In other words, while the location of \mathbf{z}^t can be arbitrary in the representation space, the change of image representation between time points $\mathbf{z}^s - \mathbf{z}^t$ is only allowed in the $\boldsymbol{\tau}$ direction.

To learn an encoder \mathbf{g} that satisfies the cosine constraint, we train a standard autoencoder that models \mathbf{g} as a neural network with parameters θ and simultaneously determines the parameters ϕ of a decoder network \mathbf{d} to reconstruct the input image from the encoded representation. In doing so, the learned representations encode all morphological information of the brain beyond the single factor of interest. To impose the cosine constraint in the autoencoder, we propose to add a cosine loss for each image pair to the standard Mean-Squared Error loss $\text{MSE}(\cdot, \cdot)$ of the autoencoder as a soft constraint, i.e.,

$$\min_{\theta, \phi, \boldsymbol{\tau}} \sum_{\mathbf{I} \in \mathcal{I}} \text{MSE}(\mathbf{I}, \mathbf{d}(\mathbf{g}(\mathbf{I}; \theta); \phi)) - \lambda \sum_{(\mathbf{I}^t, \mathbf{I}^s) \in \mathcal{S}} \cos(\mathbf{g}(\mathbf{I}^s; \theta) - \mathbf{g}(\mathbf{I}^t; \theta), \boldsymbol{\tau}), \quad (6)$$

with λ being the parameter weighting the two terms. As a result, the objective function encourages the encoder to learn the low-dimensional representation of images while encouraging the development of brain representation of all subjects to be colinear with $\boldsymbol{\tau}$.

4. Experiments

4.1. Experimental setup

Datasets We evaluated the proposed LSSL on a synthetic and two longitudinal neuroimaging datasets. The synthetic dataset contained 512 subjects. Each subject consisted of an image pair \mathbf{I}^1 and \mathbf{I}^2 , whose difference was regarded as the developmental effect over time. Either image in the pair contained four Gaussian patterns (Fig. 3(a)). The magnitude of the two diagonal Gaussians were randomly sampled from a uniform distribution $\mathcal{U}(1, 6)$ and was kept the same for the pair. The magnitude of the two off-diagonal Gaussians simulated the ‘brain age’, which was sampled from $\psi \sim \mathcal{U}(1, 4)$ in the first image and set to $\psi + \Delta\psi$ for the second image with $\Delta\psi \sim \mathcal{U}(0.1, 2)$ being the age-related increase in magnitude between the two images (Fig. 3(a)). Gaussian noise of SNR=8 was added to each image. Training on the 512 image pairs, the goal of the synthetic experiment was to show that LSSL can disentangle a direction $\boldsymbol{\tau}$ in the latent space encoding the off-diagonal developmental pattern and that the estimated brain age ψ (projections along $\boldsymbol{\tau}$) accurately correlates with the ground-truth ψ .

Next, we evaluated LSSL on the 2641 structural MRIs of 811 subjects from the Alzheimer’s Disease Neuroimaging Initiative (ADNI1²). The dataset consisted of 229 normal control subjects (age: 76 ± 5.0 years), 397 subjects diagnosed with Mild Cognitive Impairment (74.9 ± 7.4 years), and 185 subjects with Alzheimer’s Disease (75.3 ± 7.6 years). The longitudinal MRI of each subject was composed of up to 8 scans (acquired within a 4 year study period) that we were able to successfully preprocess. In line with our prior studies (Adeli et al., 2020; Zhao et al., 2020), the preprocessing consisted of denoising, bias field correction, skull stripping, affine registration to a template, re-scaling to a $64 \times 64 \times 64$ volume, and transforming image intensities within the brainmask to z-scores. We constructed 3141 image pairs based on the criteria that each pair belonged to the same subject and had at least one year interval in scan time.

Another neurological disorder known to accelerate brain aging is alcohol dependence, which can cause gradual deterioration in the gray and white matter tissue (Pfefferbaum et al., 2014; Zahr and Pfefferbaum, 2018). Therefore, the second MRI dataset was comprised of 1499 T1-weighted MRIs of 274 Normal Controls (age: 47.3 ± 17.6) and 329 patients diagnosed with alcohol dependence (age: 49.3 ± 10.5) according to the DSM-IV criteria (Bell, 1994) (referred to as the alcohol data set). 74 participants of the alcoholic group were also human immunodeficiency virus (HIV) positive. Each subject had up to 13 longitudinal scans. 1071 image pairs were constructed from this dataset based on the above one-year-interval criterion. The study was approved by the institutional review boards of Stanford University School of Medicine and SRI International. All MRIs were pre-processed using the prior pipeline.

Implementation of LSSL As LSSL is a formulation on latent representations (which is not specific to any design of the decoder \mathbf{d} and encoder \mathbf{g}), we constructed rather simplistic and standard autoencoders so that findings revealed herein are likely to generalize to more advanced autoencoder structures. For the synthetic dataset, we designed the encoder of LSSL as 3 stacks of 3×3 convolution

² Data publicly available at <http://adni.loni.usc.edu/>

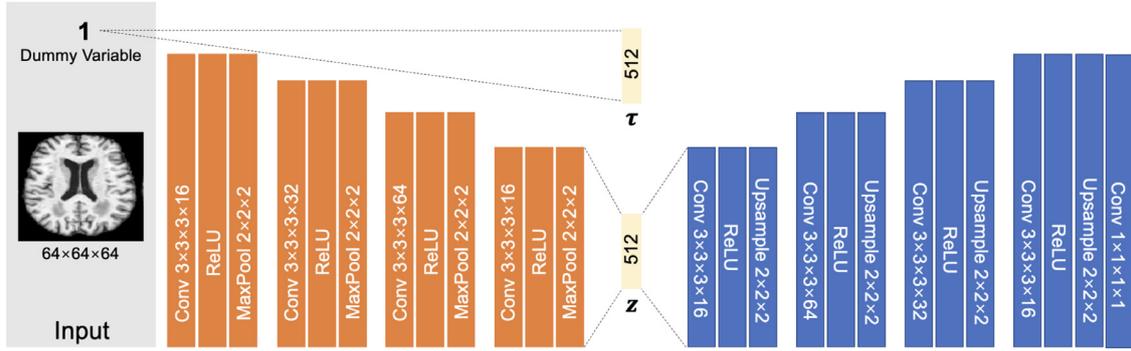


Fig. 2. Network structure of LSSL. Orange blocks correspond to the encoder and blue blocks to the decoder networks. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

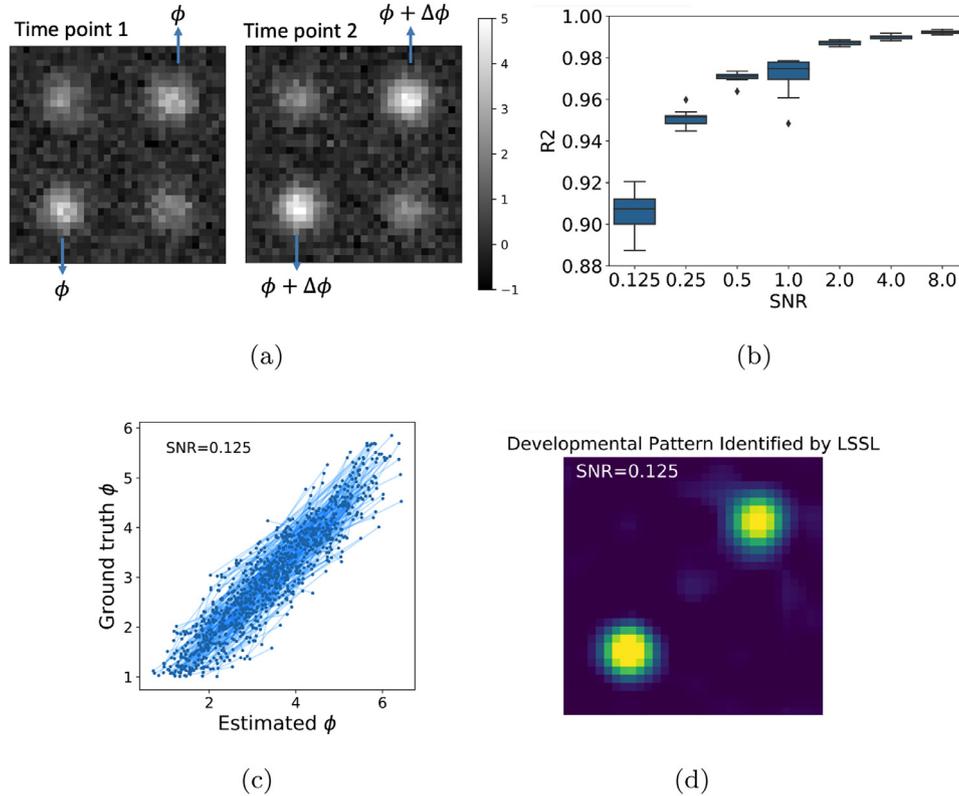


Fig. 3. (a) A synthetic image pair representing a subject's developmental effects, which were quantified by the magnitude change $\Delta\phi$ of the two off-diagonal Gaussians; (b) R2 between estimated and ground-truth ϕ with respect to noise level; (c) Group-level correlation for SNR = 0.125. Each line connects two points associated with an image pair; (d) Identified developmental patterns for SNR = 0.125.

with feature dimension $\{2, 4, 8\}$, tanh activation, and max-pooling layers, which resulted in a 16 dimensional representation. The decoder employed a reverse structure of the encoder. The unit vector τ was embedded in the network as the output of a fully connected layer applied to a dummy input scalar (Fig. 2), which made τ a global variable and independent of the specific input image. For the MRI datasets, the encoder was 4 stacks of $3 \times 3 \times 3$ convolution with dimension $\{16, 32, 64, 16\}$, ReLU activation, and max-pooling layers (Fig. 2). A fully connected layer resulted in a 512 dimensional representation space.

For each dataset, we set a non-informative hyperparameter λ as the ratio between the number of images ($|I|$) and the number of image pairs ($|S|$) to balance the number of samples in the two loss terms (Eq. (6)). This ratio resulted in a reasonable balance between the two loss components in the objective function (see Supplement Fig. S3).

We trained the models for 100 epochs using an Adam optimizer with an initial learning rate of 0.001. In the MR experiments, we reduced the learning rate by a factor of 0.2 if the training loss stopped decreasing for 5 epochs. We confirmed convergence based on the train loss curve in all experiments. The models were implemented in Keras 2.2.4 and ran on an Nvidia Quadro P6000 GPU with 2 GB memory. Each training run in the MR experiments took approximately 2–4 h.

Evaluation on the MR datasets For each dataset, we first trained LSSL on the collection of image pairs (subjects with only one MR were omitted from training) and applied the trained model to derive the brain age ψ (projections along τ) for all the MRIs in the dataset. Note, the ground-truth diagnosis labels and the chronological age were omitted from the self-supervised training. Then, the estimated brain age was correlated with chronological age in the control subjects to verify the quality of disentanglement. The brain

age of control subjects was compared to that of the diagnosed patients to reveal disease effects. Note, this type of analysis was unique to LSSL and could not be achieved by existing disentanglement approaches based on statistical formulation (e.g., Kingma and Welling, 2014; Higgins et al., 2017; van den Oord et al., 2018; Chen et al., 2020) because those methods do not estimate latent directions associated with underlying factors.

LSSL also learned a representation vector for each MRI. We evaluated the quality of these representations by using them to classify diagnosis labels of individuals in a supervised setting. The classification was evaluated by 5-fold cross-validation, where the folds were split based on subjects; i.e., images of a single subject belonged to the same fold. After splitting folds, we performed the classification using both cross-sectional (i.e., based on single time points) and longitudinal models. The cross-sectional model discarded the temporal information within each subject and treated each image as an independent sample. The classifier was designed as a Multi-Layer Perceptron containing two fully connected layers of dimension 512 and 64 with ReLU activation. The longitudinal model was a Recurrent Neural Network (RNN), whose input was the longitudinal sequence of representation vectors for each subject. The RNN mapped each representation vector within the input sequence to a 16 dimensional vector, which was fed into a single layer GRU network with 16 hidden units (Ouyang et al., 2021) to predict the diagnosis label at each visit. In a separate experiment, we fine-tuned the LSSL representation by incorporating the encoder \mathbf{g} (pre-trained by LSSL) into the classification models and then cross-validated this end-to-end classifier. Classification accuracy was measured by balanced accuracy (bAcc) to account for different number of samples in each cohort (Zhao et al., 2020).

Baselines We compared the classification accuracy and speed of convergence of the end-to-end classifier to those whose encoders were pre-trained by several other state-of-the-art representation learning methods. For fair comparison, we used the same encoder architecture for all methods, as our goal was to show the superiority of our self-supervised representation learning rather than obtaining state-of-the-art results on any of the two datasets with more complex encoder architectures. As LSSL was conceptually related to a wide range of works, we selected several representative methods from unsupervised training (AE and VAE), factor disentanglement (β -VAE Higgins et al., 2017), self-supervised learning (SimCLR Chen et al., 2020), to a longitudinal framework based on Contrastive Predictive Coding (CPC van den Oord et al., 2018). Note, the pre-training of CPC already contained an auto-regressive model on top of the encoder, which reduced the representation to ‘context’ features, i.e., a 16-D vector followed by a GRU with 16 hidden units. Therefore, the longitudinal prediction of CPC was directly based on the context features instead of the 512-D representation.

4.2. Results of synthetic experiments

We trained LSSL on the 512 image pairs and derived the brain age ψ for each synthetic image. As the scale of the coordinates in the latent space is not uniquely determined (one can rescale the latent space by rescaling the network parameters of the encoder and decoder), we normalized ψ such that its mean and standard deviation were matched to those of the ground truth. Note, this normalization was solely for intuitive interpretation of the results. Next, we quantified the correlation between the ground-truth and estimated ψ by the R2 score (Da Costa Lewis, 2005). Fig. 3b indicates that this correlation was nearly perfect ($R^2 = 0.99$) for a high SNR = 8 and remained high for a low SNR = 0.125. This global correlation in the range of $\psi \in [1, 6]$ was learned from pairs of data whose maximum difference in brain age $\Delta\psi \leq 2$ (Fig. 3c). To identify the spatial pattern associated with increasing brain age, we varied the average latent representation $\bar{\mathbf{z}}$ (across all images)

along τ and $-\tau$ by one unit and visualized the resulting difference between the two reconstructed images:

$$\mathbf{d}(\bar{\mathbf{z}} + \tau; \phi) - \mathbf{d}(\bar{\mathbf{z}} - \tau; \phi) \quad (7)$$

The resulting image (Fig. 3d) shows that LSSL accurately estimated the developmental pattern in the two off-diagonal Gaussians even for the low SNR setting.

4.3. Longitudinal study of Alzheimer’s disease

We trained LSSL on the 3141 pairs of MRIs from ADNI and then applied the model to derive the brain age ψ for all 2641 MRI in the dataset. The mean and standard deviation of the estimated ψ were normalized according to the chronological age in the dataset. Fig. 4a shows the brain age of the control subjects versus their chronological age. According to the fitting of a quadratic mixed effect model, brain age and chronological age exhibited a nearly linear relationship over the entire age span of the dataset. To ensure this correlation was not a result of model overfitting, we measured the Pearson’s correlation only on the 43 control MRIs that were not a part of the training set (see Supplement Fig. S4). The ‘global’ correlation in the range of 60 to 90 years was derived solely from the ordinal information from subject-specific image pairs (maximum 4 years part) without using the ground-truth chronological age of subjects.

Next, we qualitatively show that the estimated ψ was a more accurate marker for brain age compared to the chronological age. Fig. 5a displays the brain of 7 normal control subjects with the same chronological age (of 80 years), yet their estimated brain age ranged from 62.9 years to 92.5 years. This large variance visually relates to the difference in ventricle size and cortical thickness shown in Fig. 5a. It is also consistent with the variance of brain age plotted in Fig. 4a for any given chronological age. On the other hand, control subjects with the same estimated brain age shown in Fig. 5b had similar brain appearance despite that they had various chronological age. These results not only support the efficacy of our brain age estimation but also suggest that supervised training based on chronological age may be a flawed strategy for learning the brain age even within the control cohort.

According to Fig. 4b, the brain age of AD patients was generally higher than chronological age reflecting the neurodegenerative nature of AD that accelerated brain aging. This phenomenon can also be inferred from Fig. 4c, where we computed the ‘brain age slope’ by fitting a simple linear regression for each subject (with at least 2 images) on their brain age across visits. In doing so, we see that the control group had an average aging speed (slope) close to 1, indicating the consistency between the progression rate of brain age and of chronological age. In comparison, brain aging of the AD group was significantly faster ($p < 0.001$, two-sample t -test). Interestingly, the MCI group, representing a transitional state between control and AD, had an intermediate aging speed, which was significantly faster than normal and slower than AD ($p < 0.001$, two-sample t -tests). Moreover, we observe that the gap between brain age and chronological age was larger in younger AD patients than the older ones (Fig. 4b), which indicates the disease effect was more prominent in the younger brain. This phenomenon was quantitatively supported by fitting a linear regression between the aging speed (slope) and age in each cohort (Fig. 4d). While the AD subjects had more accelerated aging at younger ages, their aging speed was not different from that of the oldest old in the control cohort. Again, the MCI group exhibited an intermediate effect between AD and control subjects. Lastly, to ensure that these results were not specific to the resolution of input images, we repeated the above experiments based on the $80 \times 80 \times 80$ input resolution, which resulted in similar findings (Supplement Fig. S1).

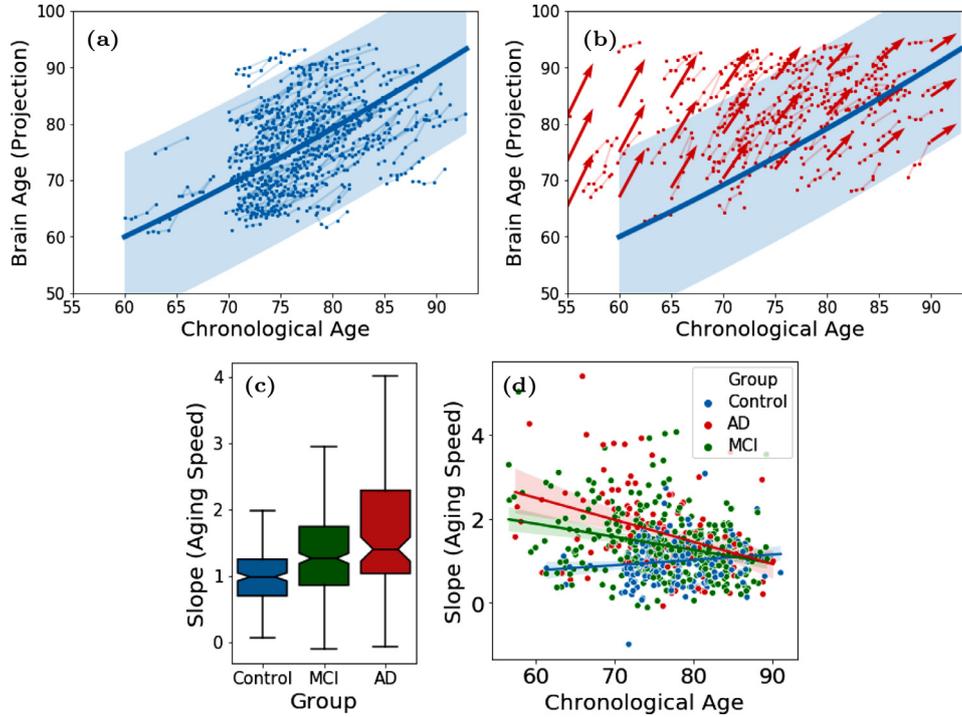


Fig. 4. (a) Brain age of 229 control subjects from ADNI1; (b) Brain age of 185 AD patients (red) overlaid with normal developmental trajectory (blue); (c) Speed of brain aging (slope of ψ over time) for the 3 diagnosis groups and (d) as a function of chronological age of the ADNI1 subjects. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

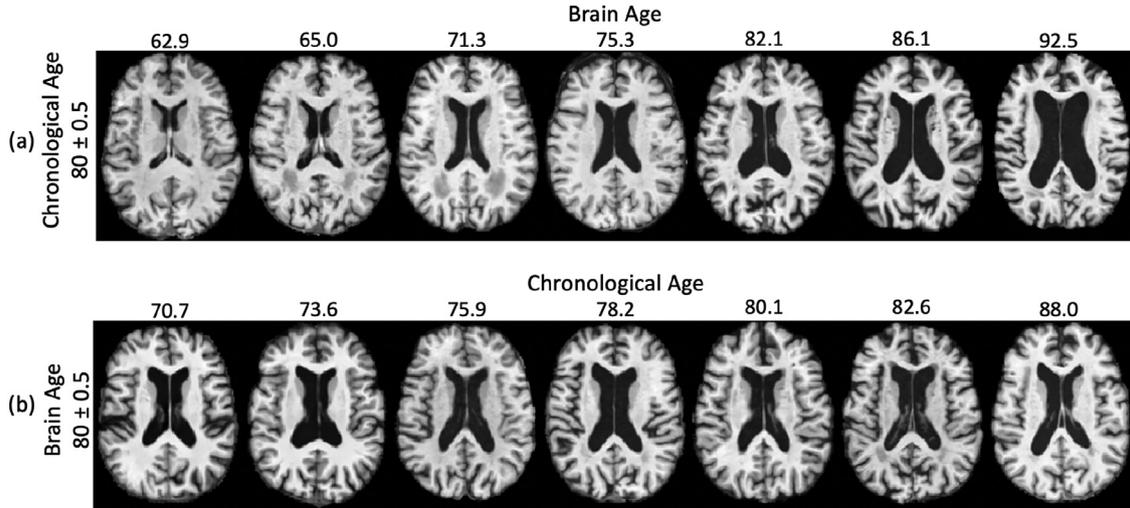


Fig. 5. (a) Control subjects with the same chronological age of 80 years had distinct brain appearance and brain age estimated by LSSL. (b) Control subjects with the same brain age of 80 years exhibited similar brain appearance.

We also assessed the quality of disentanglement by simulating the average brain at different brain ages. We constructed the average representation associated brain age ψ via

$$\psi \tau + \frac{1}{N} \sum_{i=1}^N (\mathbf{z}^i - \mathbf{z}^i \tau^\top \tau), \quad (8)$$

where the first term corresponds to the representation of brain age and the second term captures factors independent from brain age, i.e., the group average of the components orthogonal to τ . By decoding this age-dependent representation, we observe a pattern of enlargement in the ventricle and loss of brain tissues as age increases (Fig. 6). This pattern converges with the current understanding of brain aging in the neuroscience literature (Sullivan and Pfefferbaum, 2008).

Lastly, we classified AD patients from control subjects based on the learned representations. In the cross-sectional setting, the representations learned by LSSL enabled more accurate prediction than the baselines except for SimCLR (Table 1 CNN), and the bAcc of our model closely matched up to SimCLR after fine-tuning the encoder in the cross-sectional setting. On the other hand, when performing classification based on the longitudinal sequences (CNN + RNN), the bAcc associated with the LSSL representations increased and outperformed all baselines, which was also the case after fine-tuning the encoders. This accuracy improvement over cross-sectional CNN indicates that LSSL resulted in informative temporal trajectories of individuals in the representation space, which could only be learned by the RNN models. Moreover, when performing the fine-tuning in an end-to-end set-

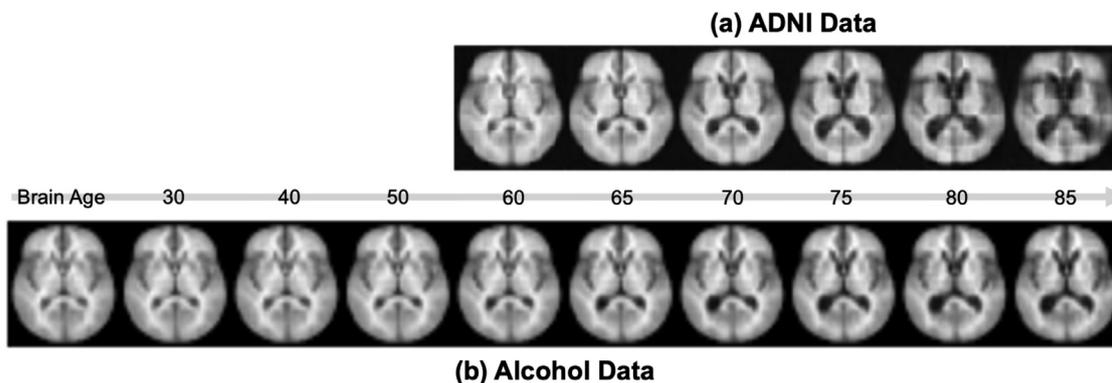


Fig. 6. Simulated average brain at different brain ages from the ADNI and alcohol data.

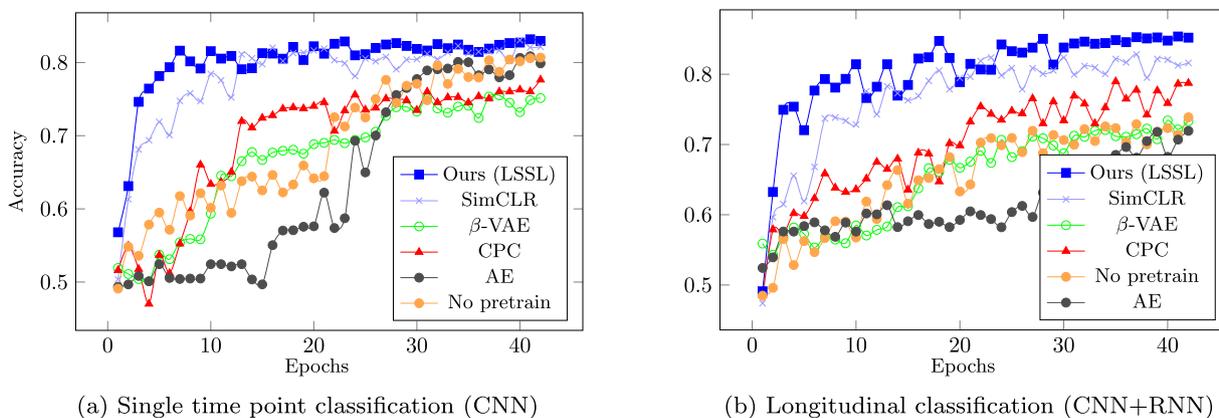


Fig. 7. ADNI Dataset - Average bAcc in the first 45 epochs over the 5 testing folds when training end-to-end classification based on pretrained encoders.

Table 1

ADNI data set: balanced accuracy (bAcc) of cross-sectional and longitudinal classification with and w/o fine-tuning the encoder. Best result in each column is in bold and the second best is underlined.

Pre-training	ADNI	
	CNN	CNN + RNN
	frozen / fine-tuned	frozen / fine-tuned
No pretrain	/ 80.6	/ 74.6
AE	58.6 / 81.7	62.1 / 71.3
VAE (Kingma and Welling, 2014)	58.9 / 75.7	62.8 / 71.9
β -VAE (Higgins et al., 2017)	56.1 / 77.2	76.3 / 78.4
SimCLR (Chen et al., 2020)	78.0 / 84.4	80.7 / 84.7
CPC (van den Oord et al., 2018)	65.5 / 78.6	66.7 / 80.4
Ours (LSSL)	<u>72.0 / 84.1</u>	81.8 / 87.0

ting, the encoder pre-trained by LSSL converged the fastest in both cross-sectional and longitudinal settings (see Fig. 7) despite that all methods started from random predictions (as the MLP layers were randomly initialized). When using the most accurate implementation (LSSL + RNN + fine-tuning), the bAcc for distinguishing MCI from controls (69.9%) and AD (69.5%) were also higher than those without pretraining of LSSL (68.3% and 66.5%). Note, the classification accuracies are in line with the literature (Oh et al., 2019), which ranges from 60% to 75% bAcc depending on the progressiveness of the MCI subjects used for classification.

4.4. Longitudinal study of alcohol dependence

Similar to the previous experiment, we first trained LSSL to estimate the brain age for all MR images and then visualized the brain age of the control subjects (Fig. 8a). To put the results in line with

the ADNI experiments, we normalized the projection ψ with respect to the age range of the ADNI dataset (i.e., confined to 60 to 90 years and then applied to the entire age range of the alcohol data set, Fig. 8a). Compared to the approximately linear aging pattern between age 60 to 90, the aging of the control subjects exhibited a quadratic pattern over a longer life span, where the aging speed was slower for younger subjects (e.g., between 20 and 60 years) compared to the older subjects (after 60 years, Fig. 8a). Similar to the ADNI experiment, the alcoholics also exhibited higher brain age than normal controls (Fig. 8b), which was supported by their slopes (aging speed) being significant larger than normal ($p < 0.001$, Fig. 8c). Different from the ADNI experiment was that the aging speed of alcohol-dependent subjects was always faster than normal controls regardless of their chronological age (Fig. 8d). Furthermore, older subjects had a larger gap between brain and chronological age (Fig. 8b), which indicates an accumulative alcohol effect. The accumulative alcohol effect is also supported by the observation that chronic drinking gradually deteriorates brain structure resulting in more severe alcohol effect in older subjects (Zhao et al., 2019b).

In Fig. 6, we simulated images of different brain ages for the alcohol dataset. We observe that the simulated brains closely mimic the ones from the ADNI experiment from age 60 to 70, a range where the two datasets highly overlap. However, the simulated brains from the alcohol experiment showed less pronounced aging effect after age 70 compared to the ADNI results. This was potentially due to the few older subjects in the alcohol dataset compared with ADNI, so the model conservatively extrapolated the aging pattern for the older age range.

Prior literature indicates that alcohol dependence is only weakly separable from the control group (Adeli et al., 2018;

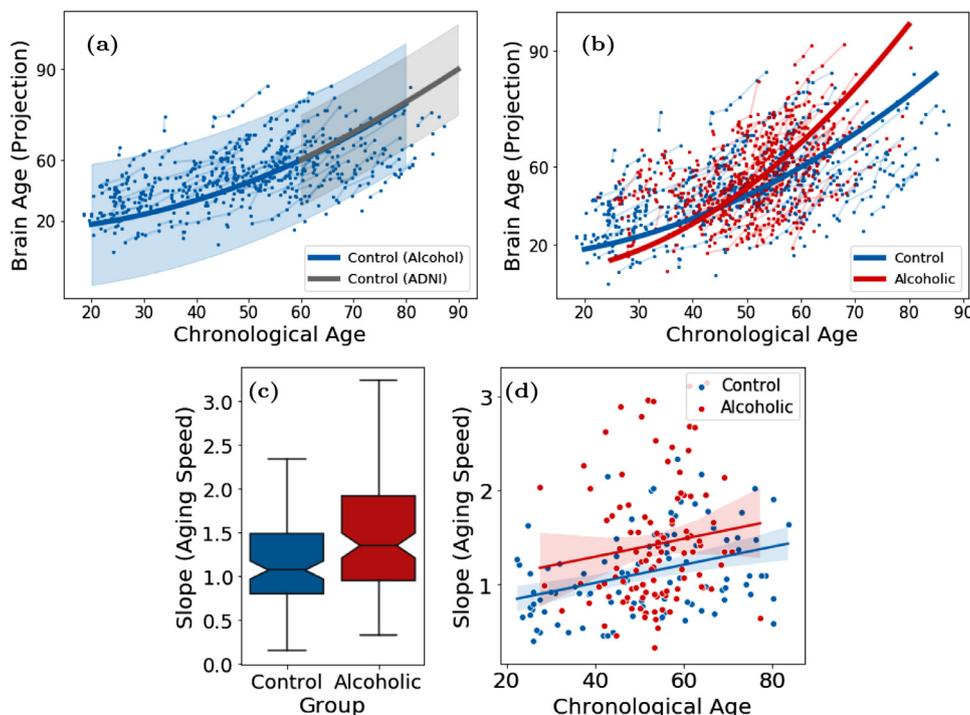


Fig. 8. (a) Brain age of 274 normal control subjects from the alcohol dataset (blue) overlaid with the ADNI1 controls (gray); (b) Brain age of 329 alcohol-dependent subjects (red); (c,d) Speed of brain aging for all subjects with two or more images in the alcohol data set. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2

Alcohol dataset: balanced accuracy (bAcc) of cross-sectional and longitudinal classification with and w/o fine-tuning the encoder. Best result in each column has bold typeset, and the second best is underlined.

Pre-training	Alcohol dependence	
	CNN	CNN + RNN
Model	frozen/fine-tuned	frozen/fine-tuned
No pretrain	/ 69.5	/ 52.8
AE	58.8 / 69.1	52.1 / 53.2
VAE (Kingma and Welling, 2014)	55.4 / 70.2	63.0 / 65.6
β -VAE (Higgins et al., 2017)	52.1 / 67.5	60.8 / 61.0
SimCLR (Chen et al., 2020)	63.2 / <u>68.7</u>	<u>66.2</u> / <u>69.3</u>
CPC (van den Oord et al., 2018)	51.9 / 67.5	62.0 / 63.2
Ours (LSSL)	<u>62.9</u> / 71.7	67.0 / 72.0

Ouyang et al., 2021), which is echoed in our results by the significantly lower bAcc in the frozen setting (all methods < 65 %, Table 2) compared to the ADNI experiment. Nevertheless, all accuracy scores associated with LSSL were significantly better than chance based on the Fisher’s exact test ($p < 0.001$). Moreover, LSSL resulted in the fastest converging rate and highest accuracy upon convergence in the fine-tuning setting compared to the baselines (Fig. 9). The more challenging task of classifying alcohol dependence compared to the AD classification also took the RNN longer to converge for several baselines, which was not the case for LSSL.

5. Discussion

In recent years, many studies have used supervised models to estimate brain age from structural MRIs. These models have to be trained on a healthy population to predict chronological age to establish the normal association between brain age and brain structures (Kaufmann et al., 2019; Smith et al., 2020; Zhao et al., 2019). Then the trained model is applied to a diseased cohort to examine the brain age gap (difference between the estimated brain age

and chronological age) induced by the disease (Kaufmann et al., 2019). This gap, however, is likely to be biased by the ‘domain shift’ between the training and testing sets. LSSL alleviates this issue by training on both control and diseased subjects (without using their labels), which results in an analysis impartial to cohort bias. Moreover, the chronological age is a sub-optimal ground-truth for brain age in supervised models as there are multiple modes of brain aging within the healthy population due to genetic influence (Smith et al., 2020) (see also Fig. 5). LSSL addresses this challenge by omitting the supervision of chronological age and only using ordinal information of within-subject scans, which permits the characterization of heterogeneity across subjects.

The advantage of LSSL over other unsupervised/self-supervised baselines is evident from our post-hoc classification, which revealed LSSL could learn discriminative cues within the representations even without using the group labels for training. However, the AD classification of LSSL did not rival with the highest accuracy score reported on the ADNI dataset (Liu et al., 2017; 2019), which was expected as we refrained from extensively exploring network architectures for the encoder. This type of exploration is a research direction orthogonal to the proposed self-supervised learning strategy in the representation space. This self-supervised learning strategy outperformed the baselines, all of which used the same encoder setting. The setting was based on the most basic components used in standard CNNs (convolution, ReLU, and max-pooling). Therefore, we expect the findings revealed herein are likely to generalize to more advanced encoder architectures.

Limitation As longitudinal studies are typically designed to examine the influence from time-dependent factors, we did not model Condition 2, i.e., the independence between ψ and other static factors, such as gender. This theoretically makes LSSL only result in a necessary condition for disentanglement. In practice, we can examine whether Condition 2 holds via post-hoc analyses. For example, in both experiments, the speed of brain aging was not significantly different between males and females ($p > 0.05$ two-

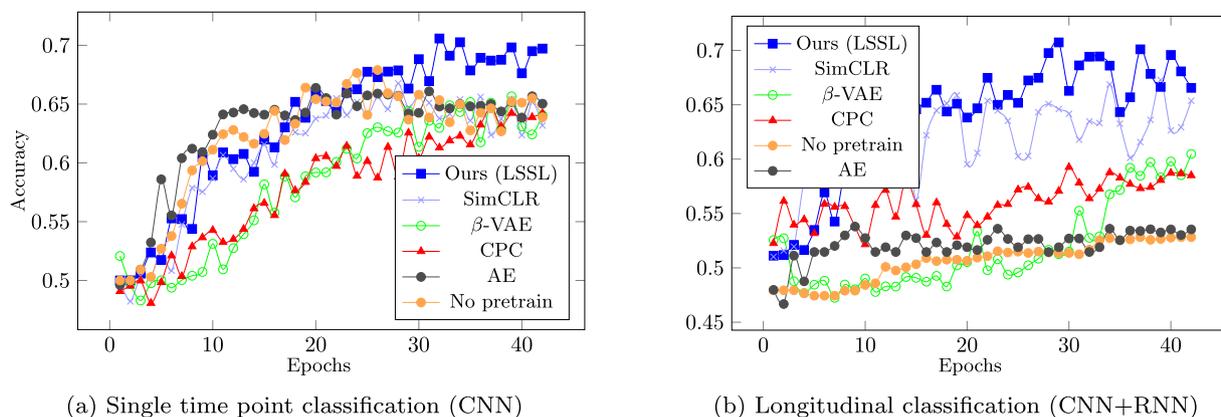


Fig. 9. Average accuracy in the first 45 epochs for classifying alcohol dependence.

sided two-sample t -test) indicating an exact disentanglement between brain age and gender.

Another limitation of LSSL is that it does not eliminate the possible confound by other time-evolving factors co-occurring with brain age. For example, the change of body mass index between visits can also alter the overall brain volume and thereby the image representation (Ward et al., 2005). Although our pipeline could reduce the impact of the change in brain volume by affinely registering each scan of the longitudinal MR images to a template, a systematic way of modeling confounders during the training of LSSL needs to be further explored.

Finally, the current formulation of LSSL can only disentangle one time-variant factor related to brain age. This constraint limited our analysis to focus on abnormalities in brain age due to a neurological condition (such as AD). To model other disease effects not related to accelerated aging, a future research direction of LSSL is to jointly disentangle two orthogonal directions in the latent space so that one could separately characterize disease progression and brain aging.

6. Conclusion

In this work, we proposed a self-supervised representation learning framework called LSSL that incorporated theoretical advantages from the repeated measures design in longitudinal neuroimaging studies. The explicit longitudinal self-supervision permitted separate definitions for the factor and representation spaces, thereby omitting the ambiguity often encountered in fully unsupervised disentanglement models. Based on optimizing the colinearity between a global direction in the representation space and the developmental trajectories from subject-specific image pairs, LSSL successfully disentangled the factor of brain aging in the representation space, which was used to characterize normal aging pattern across the life span and to reveal the accelerated aging effects of Alzheimer's Disease and alcohol dependence. Compared to several other state-of-the-art representation learning methods, the pre-trained encoder and representations learned by LSSL are more suitable for supervised classification of diagnosis labels in various settings, indicated by faster convergence and higher (or equally high) prediction accuracy upon convergence.

Declaration of Competing Interest

Authors declare that they have no conflict of interest.

CRediT authorship contribution statement

Qingyu Zhao: Conceptualization, Methodology, Software, Formal analysis, Writing - original draft. **Zixuan Liu:** Methodology, Software, Formal analysis. **Ehsan Adeli:** Conceptualization, Methodology, Writing - review & editing. **Kilian M. Pohl:** Data curation, Writing - review & editing, Supervision.

Acknowledgments

This work was supported by NIH Grants [MH113406](#), [AA005965](#), [AA010723](#), and [AA017437](#), and by Stanford HAI AWS Cloud Credit.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.media.2021.102051](#)

References

- Adeli, E., Kwon, D., Pohl, K., 2018. Multi-label transduction for identifying disease comorbidity patterns. In: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part III, pp. 575–583. doi:[10.1007/978-3-030-00931-1_66](#).
- Adeli, E., Zhao, Q., Zahr, N.M., Goldstone, A., Pfefferbaum, A., Sullivan, E.V., Pohl, K.M., 2020. Deep learning identifies morphological determinants of sex differences in the pre-adolescent brain. *NeuroImage* 223, 1–13.
- Aghili, M., Tabarestani, S., Adjouadi, M., Adeli, E., 2018. Predictive modeling of longitudinal data for Alzheimer's disease diagnosis using RNNs. In: International Workshop on Predictive Intelligence In Medicine. Springer, pp. 112–119.
- Bell, C.C., 1994. DSM-IV: diagnostic and statistical manual of mental disorders. *JAMA* 272 (10), 828–829.
- Bernal-Rusiel, J.L., Greve, D.N., Reuter, M., Fischl, B., Sabuncu, M.R., ADNI, 2013. Statistical analysis of longitudinal neuroimage data with linear mixed effects models. *NeuroImage* 66, 249–260.
- Caruana, E., Roman, M., Hernández-Sánchez, J., Solli, P., 2015. Longitudinal studies. *J. Thorac. Dis.* 7 (11), 537–540. doi:[10.3978/j.issn.2072-1439.2015.10.63](#).
- Chen, M., Shi, X., Zhang, Y., Wu, D., Guizani, M., 2017. Deep features learning for medical image analysis with convolutional autoencoder neural network. *IEEE Trans. Big Data* PP, 1. doi:[10.1109/TBDATA.2017.2717439](#).
- Chen, R.T.Q., Li, X., Grosse, R.B., Duvenaud, D.K., 2018. Isolating sources of disentanglement in variational autoencoders. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc., pp. 2610–2620.
- Cash, P., Stanković, T., Štorga, M., 2016. An Introduction to Experimental Design Research. In: Cash, P., Stanković, T., Štorga, M. (Eds.), *Experimental Design Research*. Springer, Cham, pp. 3–12.
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E., 2020. A simple framework for contrastive learning of visual representations. In: *International Conference on Machine Learning*, pp. 1597–1607.
- Cole, J., Poudel, R., Tsagkrasoulis, D., Caan, M., Steves, C., Spector, T., Montana, G., 2016. 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](#).
- Cui, R., Liu, M., Alzheimer's Disease Neuroimaging Initiative, et al., 2019. RNN-based longitudinal analysis for diagnosis of Alzheimer's disease. *Comput. Med. Imaging Graph.* 73, 1–10.

- Da Costa Lewis, N., 2005. A Primer in Applied Regression Analysis. In: Energy Risk Modeling. Finance and Capital Markets Series. Palgrave Macmillan, London, pp. 151–168.
- Dennis, N., Cabeza, R., 2008. Neuroimaging of healthy cognitive aging. In: The Handbook of Aging and Cognition, pp. 1–54.
- Doersch, C., Zisserman, A., 2017. Multi-task self-supervised visual learning. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2051–2060.
- Dong, C., Loy, C.C., He, K., Tang, X., 2014. Learning a deep convolutional network for image super-resolution. In: European Conference on Computer Vision. Springer, pp. 184–199.
- Elliott, M., Belsky, D., Knodt, A., Ireland, D., Melzer, T., Poulton, R., Ramrakha, S., Caspi, A., Moffitt, T., Hariri, A., 2019. Brain-age in midlife is associated with accelerated biological aging and cognitive decline in a longitudinal birth cohort. *Mol. Psychiatry* 1–10. doi:10.1038/s41380-019-0626-7.
- Franke, K., Gaser, C., 2019. Ten years of brainage as a neuroimaging biomarker of brain aging: what insights have we gained? *Front. Neurol.* 10, 1–26. doi:10.3389/fneur.2019.00789.
- Ghazi, M.M., Nielsen, M., Pai, A., Cardoso, M.J., Modat, M., Ourselin, S., Sørensen, L., Alzheimer's Disease Neuroimaging Initiative, et al., 2019. Training recurrent neural networks robust to incomplete data: application to Alzheimer's disease progression modeling. *Med. Image Anal.* 53, 39–46.
- Habeck, C., Stern, Y., 2010. Multivariate data analysis for neuroimaging data: overview and application to Alzheimer's disease. *Cell Biochem. Biophys.* 58, 53–67. doi:10.1007/s12013-010-9093-0.
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R., 2019. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*.
- Higgins, I., Amos, D., Pfau, D., Racaniere, S., Matthey, L., Rezende, D., Lerchner, A., 2018. Towards a definition of disentangled representations. *arXiv:1812.02230*.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M.M., Mohamed, S., Lerchner, A., 2017. beta-VAE: learning basic visual concepts with a constrained variational framework. *ICLR*.
- Jack, C., Bernstein, M., Fox, N., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P., Whitwell, J., Ward, C., Dale, A., Felmlee, J., Gunter, J., Hill, D., Killiany, R., Shuff, N., Fox-Bosetti, S., Lin, C., Studholme, C., Weiner, M., 2008. The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods. *J. Magn. Reson. Imaging* 27, 685–691. doi:10.1002/jmri.21049.
- Jing, L., Tian, Y., 2019. Self-supervised visual feature learning with deep neural networks: a survey. *arXiv:1902.06162*
- Johnson, J., Alahi, A., Fei-Fei, L., 2016. Perceptual losses for real-time style transfer and super-resolution. In: European Conference on Computer Vision. Springer, pp. 694–711.
- Kaufmann, T., Meer, D., Doan, N., Schwarz, E., Lund, M., Agartz, I., Alnäs, D., Barch, D., Baur-Streibl, R., Bertolino, A., Bettella, F., Beyer, M., Båen, E., Borgwardt, S., Brandt, C., Buitelaar, J., Celius, E., Cervenka, S., Conzelmann, A., Westlye, L., 2019. Common brain disorders are associated with heritable patterns of apparent aging of the brain. *Nat. Neurosci.* 22, 1617–1623. doi:10.1038/s41593-019-0471-7.
- Kim, H., Mnih, A., 2018. Disentangling by factorising. In: Dy, J., Krause, A. (Eds.), Proceedings of the 35th International Conference on Machine Learning. PMLR, Stockholm Sweden, pp. 2649–2658.
- Kingma, D.P., Welling, M., 2014. Auto-encoding variational bayes. In: 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings.
- Kolesnikov, A., Zhai, X., Beyer, L., 2019. Revisiting self-supervised visual representation learning. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1920–1929.
- Larsson, G., Maire, M., Shakhnarovich, G., 2017. Colorization as a proxy task for visual understanding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6874–6883.
- Lipton, Z. C., Kale, D. C., Elkan, C., Wetzel, R., 2015. Learning to diagnose with LSTM recurrent neural networks. *arXiv preprint arXiv:1511.03677*
- Liu, C., Cripe, T., Kim, M.-O., 2010. Statistical issues in longitudinal data analysis for treatment efficacy studies in the biomedical sciences. *Mol. Therapy* 18, 1724–1730. doi:10.1038/mt.2010.127.
- Liu, M., Li, F., Yan, H., Wang, K., Ma, Y., Shen, L., Xu, M., 2019. A multi-model deep convolutional neural network for automatic hippocampus segmentation and classification in Alzheimer's disease. *NeuroImage* 208, 1–15. doi:10.1016/j.neuroimage.2019.116459.
- Liu, M., Zhang, J., Adeli, E., 2017. Landmark-based deep multi-instance learning for brain disease diagnosis. *Med. Image Anal.* 43, 157–168. doi:10.1016/j.media.2017.10.005.
- Locatello, F., Bauer, S., Lucic, M., Ratsch, G., Gelly, S., Schölkopf, B., Bachem, O., 2018. Challenging common assumptions in the unsupervised learning of disentangled representations. *ICML*.
- Louis, M., Couronné, R., Koval, I., Charlier, B., Durrleman, S., 2019. Riemannian geometry learning for disease progression modelling. In: Information Processing in Medical Imaging, pp. 542–553. doi:10.1007/978-3-030-20351-1_42.
- Noroozi, M., Favaro, P., 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In: European Conference on Computer Vision. Springer, pp. 69–84.
- Oh, K., Chung, Y.-C., Kim, K., Kim, W.-S., Oh, I.-S., 2019. Classification and visualization of Alzheimer's disease using volumetric convolutional neural network and transfer learning. *Sci. Rep.* 9, 1–16. doi:10.1038/s41598-019-54548-6.
- van den Oord, A., Li, Y., Vinyals, O., 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*
- Ouyang, J., Zhao, Q., Sullivan, E.V., Pfefferbaum, A., Tapert, S.F., Adeli, E., Pohl, K.M., 2021. Longitudinal pooling & consistency regularization to model disease progression from MRIs. *IEEE J. Biomed. Health Inform.* 25 (6), 1–11. doi:10.1109/JBHI.2020.3042447.
- Pfefferbaum, A., Rosenbloom, M., Chu, W., Sassoun, S., Rohlfing, T., Pohl, K., Zahr, N., Sullivan, E., 2014. White matter microstructural recovery with abstinence and decline with relapse in alcohol dependence interacts with normal ageing: a controlled longitudinal DTI study. *Lancet Psychiatry* 1, 202–212. doi:10.1016/S2215-0366(14)70301-3.
- Pillai, S., Ambrus, R., Gaidon, A., 2019. Superdepth: self-supervised, super-resolved monocular depth estimation. In: 2019 International Conference on Robotics and Automation (ICRA), IEEE, pp. 9250–9256.
- Ren, Z., Lee, Y.J., 2018. Cross-domain self-supervised multi-task feature learning using synthetic imagery. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 762–771.
- Rolinek, M., Zietlow, D., Martius, G., 2018. Variational autoencoders pursue PCA directions (by accident). *arXiv:1812.06775*
- Sabokrou, M., Khalooei, M., Adeli, E., 2019. Self-supervised representation learning via neighborhood-relational encoding. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 8010–8019.
- Santeramo, R., Withey, S., Montana, G., 2018. Longitudinal detection of radiological abnormalities with time-modulated LSTM. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. Springer, pp. 326–333.
- Smith, S., Elliott, L., Alfaro-Almagro, F., McCarthy, P., Nichols, T., Douaud, G., Miller, K., 2020. Brain aging comprises many modes of structural and functional change with distinct genetic and biophysical associations. *eLife* 9, 1–28. doi:10.7554/eLife.52677.
- Steffener, J., Habeck, C., O'Shea, D., Razlighi, Q., Bherer, L., Stern, Y., 2016. Differences between chronological and brain age are related to education and self-reported physical activity. *Neurobiol. Aging* 40, 138–144. doi:10.1016/j.neurobiolaging.2016.01.014.
- Sullivan, E., Pfefferbaum, A., 2008. Neuroradiological characterization of normal adult aging. *Br. J. Radiol.* 80 (Spec No 2), S99–108. doi:10.1259/bjr/22893432.
- Tschannen, M., Bachem, O., Lucic, M., 2018. Recent advances in autoencoder-based representation learning. *The 3rd Workshop on Bayesian Deep Learning, NeurIPS 2018*.
- Ward, M., Carlsson, C., Trivedi, M., Sager, M., Johnson, S., 2005. The effect of body mass index on global brain volume in middle-aged adults: a cross sectional study. *BMC Neurol.* 5, 1–7. doi:10.1186/1471-2377-5-23.
- Xiong, Y., Kim, H., Mehta, R., Johnson, S., Singh, V., 2019. On training deep 3D CNN models with dependent samples in neuroimaging. In: Information Processing in Medical Imaging, pp. 99–111. doi:10.1007/978-3-030-20351-1_8.
- Zahr, N.M., Pfefferbaum, A., 2018. Alcohol's effects on the brain: Neuroimaging results in humans and animal models. *Alcohol Res.* 38 (2), 183–206.
- Zhang, R., Isola, P., Efros, A.A., 2016. Colorful image colorization. In: European Conference on Computer Vision. Springer, pp. 649–666.
- Zhao, Q., Adeli, E., Honnorat, N., Leng, T., Pohl, K.M., 2019. Variational autoencoder for regression: application to brain aging analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, 11765. Springer, Cham, pp. 823–831.
- Zhao, Q., Adeli, E., Pohl, K.M., 2020. Training confounder-free deep learning models for medical applications. *Nat. Commun.* 11 (6010), 1–9.
- Zhao, Q., Pfefferbaum, A., Podhajsky, S., Pohl, K., Sullivan, E., 2019. Accelerated aging and motor control deficits are related to regional deformation of central cerebellar white matter in alcohol use disorder. *Addict. Biol.* 25 (3), 1–12. doi:10.1111/adb.12746.
- Zhao, S., Song, J., Ermon, S., 2017. InfoVAE: information maximizing variational autoencoders. *arXiv:1706.02262*.