



Self-supervised Longitudinal Neighbourhood Embedding

Jiahong Ouyang¹(✉), Qingyu Zhao¹, Ehsan Adeli¹, Edith V. Sullivan¹,
Adolf Pfefferbaum^{1,2}, Greg Zaharchuk¹, and Kilian M. Pohl^{1,2}

¹ Stanford University, Stanford, CA 94305, USA
jiahongo@stanford.edu

² SRI International, Menlo Park, CA 94025, USA

Abstract. Longitudinal MRIs are often used to capture the gradual deterioration of brain structure and function caused by aging or neurological diseases. Analyzing this data via machine learning generally requires a large number of ground-truth labels, which are often missing or expensive to obtain. Reducing the need for labels, we propose a self-supervised strategy for representation learning named Longitudinal Neighborhood Embedding (LNE). Motivated by concepts in contrastive learning, LNE explicitly models the similarity between trajectory vectors across different subjects. We do so by building a graph in each training iteration defining neighborhoods in the latent space so that the progression direction of a subject follows the direction of its neighbors. This results in a smooth trajectory field that captures the global morphological change of the brain while maintaining the local continuity. We apply LNE to longitudinal T1w MRIs of two neuroimaging studies: a dataset composed of 274 healthy subjects, and Alzheimer’s Disease Neuroimaging Initiative (ADNI, $N = 632$). The visualization of the smooth trajectory vector field and superior performance on downstream tasks demonstrate the strength of the proposed method over existing self-supervised methods in extracting information associated with normal aging and in revealing the impact of neurodegenerative disorders. The code is available at <https://github.com/ouyangjiahong/longitudinal-neighbourhood-embedding>.

1 Introduction

Although longitudinal MRIs enable noninvasive tracking of the gradual effect of neurological diseases and environmental influences on the brain over time [23], the analysis is complicated by the complex covariance structure characterizing a mixture of time-varying and static effects across visits [8]. Therefore, training deep learning models on longitudinal data typically requires a large amount of samples with accurate ground-truth labels, which are often expensive or infeasible to acquire for some neuroimaging applications [4].

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-87196-3_8) contains supplementary material, which is available to authorized users.

Recent studies suggest that the issue of inadequate labels can be alleviated by self-supervised learning, the aim of which is to automatically learn representations by training on pretext tasks (i.e., tasks that do not require labels) before solving the supervised downstream tasks [14]. State-of-the-art self-supervised models are largely based on contrastive learning [5, 10, 17, 19, 21], i.e., learning representations by teaching models the difference and similarity of samples. For example, prior studies have generated or identified similar or dissimilar sample pairs (also referred to as positive and negative pairs) based on data augmentation [6], multi-view analysis [21], and organizing samples in a lookup dictionary [11]. Enforcing such an across-sample relationship in the learning process can then lead to more robust high-level representations for downstream tasks [16].

Despite the promising results of contrastive learning on cross-sectional data [2, 7, 15], the successful application of these concepts to longitudinal neuroimaging data still remains unclear. In this work, we propose a self-supervised learning model for longitudinal data by exploring the similarity between ‘trajectories’. Specifically, the longitudinal MRIs of a subject acquired at multiple visits characterize gradual aging and disease progression of the brain over time, which manifests a temporal progression trajectory when projected to the latent space. Subjects with similar brain appearances are likely to exhibit similar aging trajectories. As a result, the trajectories from a cohort should collectively form a smooth trajectory field that characterizes the morphological change of the brain development over time. We hypothesize that regularizing such smoothness in a self-supervised fashion can result in a more informative latent space representation, thereby facilitating further analysis of healthy brain aging and effects of neurodegenerative diseases.

To achieve the smooth trajectory field, we build a dynamic graph in each training iteration to define a neighborhood in the latent space for each subject. The graph then connects nearby subjects and enforces their progression directions to be maximally aligned (Fig. 1). As such, the resulting latent space captures the global complexity of the progression while maintaining the local continuity of the nearby trajectory vectors. We name the trajectory vectors learned from the neighborhood as *Longitudinal Neighbourhood Embedding (LNE)*.

We evaluate our method on two longitudinal structural MRI datasets: one consists of 274 healthy subjects with the age ranging from 20 to 90, and the second is composed of 632 subjects from ADNI to analyze the progression trajectory of Normal Control (NC), static Mild Cognitive Impairment (sMCI), progressive Mild Cognitive Impairment (pMCI), and Alzheimer’s Disease (AD). On these datasets, the visualization of the latent space in a 2D space confirms that the smooth trajectory vector field learned by the proposed method encodes continuous variation with respect to brain aging. When evaluated on downstream tasks, we obtain higher squared-correlation (R²) in age regression and better balanced accuracy (BACC) in ADNI classifications using our pre-trained model compared to alternative self-supervised or unsupervised pre-trained models.

2 Method

We now describe LNE, a method that smooths trajectories in the latent space by local neighborhood embedding. While trajectory regularization has been explored in 2D or 3D spaces (e.g., pedestrian trajectory [1,9] and in non-rigid registration [3]), there are several challenges in the context of longitudinal MRI analysis: (1) each trajectory is measured on sparse and asynchronous (e.g., not aligned by age or visit time) time points; (2) the trajectories live in a high-dimensional space rather than a regular 2D or 3D grid space; (3) the latent representations are defined in a variant latent space that is iteratively updated. To resolve these challenges, we first propose a strategy to train based on pairwise data and translate the trajectory-regularization problem to the estimation of a smooth vector field, which is then solved by longitudinal neighbourhood embedding on dynamic graphs.

Pairwise Training Strategy. As shown in Fig. 1, each subject is associated with a trajectory (blue vectors) across multiple visits (≥ 2) in the latent space. To overcome the problem of the small number of time points in each trajectory, we propose to discretize a trajectory into multiple vectors defined by pairs of images. Compared to using the whole trajectory of sequential images as a training sample as typically done by recurrent neural networks [18], this pairwise strategy substantially increases the number of training samples. To formalize this operation, let \mathcal{X} be the collection of all MR images and \mathcal{S} be the set of subject-specific image pairs; i.e., \mathcal{S} contains all (x^t, x^s) that are from the same subject with x^t scanned before x^s . These image pairs are then the input to the Encoder-Decoder structure shown in Fig. 1. The latent representations generated by the encoder are denoted by $z^t = F(x^t)$, $z^s = F(x^s)$, where F is the encoder. Then, $\Delta z^{(t,s)} = (z^s - z^t) / \Delta t^{(t,s)}$ is formulated as the normalized trajectory vector, where $\Delta t^{(t,s)}$ is the time interval between the two scans. All $\Delta z^{(t,s)}$ in the cohort define the trajectory vector field. The latent representations are then used to reconstruct the input images by the decoder H , i.e., $\tilde{x}^t = H(z^t)$, $\tilde{x}^s = H(z^s)$.

Longitudinal Neighbourhood Embedding. Inspired by social pooling in pedestrian trajectory prediction [1,9], we model the similarity between each subject-specific trajectory vector with those from its neighbourhood to enforce the smoothness of the entire vector field. As the high-dimensional latent space cannot be defined by a fixed regular grid (e.g., a 2D image grid space), we propose to define the neighbourhood by building a directed graph \mathcal{G} in each training iteration for the variant latent space that is iteratively updated. The position of each node is defined by the starting point z^t of the vector Δz and the value (representation) of that node is Δz itself. For each node i , Euclidean distances to other nodes $j \neq i$ are computed by $P_{i,j} = \|z_i^t - z_j^t\|_2$ while the N_{nb} closest nodes of node i form its 1-hop neighbourhood \mathcal{N}_i with edges connected to i . The adjacency matrix A for \mathcal{G} is then defined as:

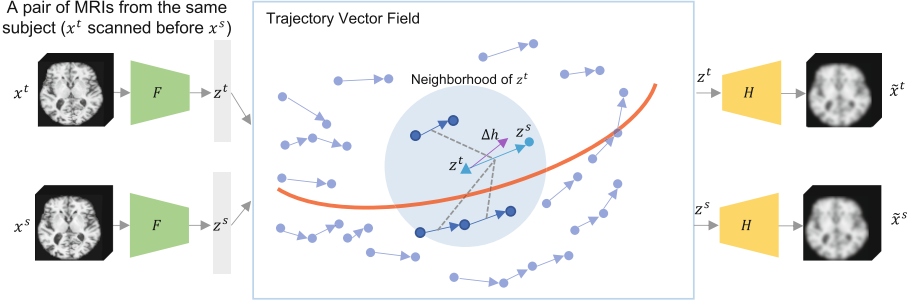


Fig. 1. Overview of the proposed method: an encoder projects a subject-specific image pair (x^t, x^s) into the latent space resulting in a trajectory vector (cyan). We encourage the direction of this vector to be consistent with Δh (purple), a vector pooled from the neighborhood of z^t (blue circle). As a result, the latent space encodes the global morphological change linked to aging (red curve). (Color figure online)

$$A_{i,j} := \begin{cases} \exp(-\frac{P_{i,j}^2}{2\sigma_i^2}) & j \in \mathcal{N}_i \\ 0, & j \notin \mathcal{N}_i \end{cases}.$$

$$\text{with } \sigma_i := \max(P_{i,j \in \mathcal{N}_i}) - \min(P_{i,j \in \mathcal{N}_i})$$

Next, we aim to impose a smoothness regularization on this graph-valued vector field. Motivated by the graph diffusion process [13], we regularize each node’s representation by a *longitudinal neighbourhood embedding* Δh ‘pooled’ from the neighbours’ representations. For node i , the neighbourhood embedding can be computed by:

$$\Delta h_i := \sum_{j \in \mathcal{N}_i} A_{i,j} D_{i,j}^{-1} \Delta z_j,$$

where D is the ‘out-degree matrix’ of graph \mathcal{G} , a diagonal matrix that describes the sum of the weights for outgoing edges at each node. As shown in Fig. 1, the blue circle illustrates the above operation of learning the neighbourhood embedding that is shown by the purple arrow.

Objective Function. As shown in [24], the speed of brain aging is already highly heterogeneous within a healthy population, and subjects with neurodegenerative diseases may exhibit accelerated aging. Therefore, instead of replacing Δz with Δh , we define $\theta_{\langle \Delta z, \Delta h \rangle}$ as the angle between Δz and Δh , and only encourage $\cos(\theta_{\langle \Delta z, \Delta h \rangle}) = 1$, i.e., a zero-angle between the subject-specific trajectory vector and the pooled trajectory vector that represents the local progression direction. As such, it enables the latent representations to model the complexity of the global progression trajectory as well as the consistency of the local trajectory vector field. To impose the direction constraint in the autoencoder, we propose to add this cosine loss for each image pair to the standard mean squared error loss, i.e.,

$$L := \mathbf{E}_{(x^t, x^s) \sim \mathcal{S}} \left(\|x^t - \tilde{x}^t\|_2^2 + \|x^s - \tilde{x}^s\|_2^2 - \lambda \cdot \cos(\theta_{\langle \Delta z, \Delta h \rangle}) \right),$$

with λ being the weighing parameter and \mathbf{E} define the expected value. The objective function encourages the low-dimensional representation of the images to be informative while maintaining a smooth progression trajectory field in the latent space. As the cosine loss is only locally imposed, the global trajectory field can be non-linear, which relaxes the strong assumption in prior studies (e.g., LSSL [24]) that aging must define a globally linear direction in the latent space. Note, our method can be regarded as a contrastive self-supervised method. For each node, the samples in its neighbourhood serve as positive pairs with the cosine loss being the corresponding contrastive loss.

3 Experiments

Dataset. To show that LNE can successfully disentangle meaningful aging information in the latent space, we first evaluated the proposed method on predicting age from 582 MRIs of 274 healthy individuals with the age ranging from 20 to 90. Each subject had 1 to 13 scans with an average of 2.3 scans spanning an average time interval of 3.8 years. The second data set comprised 2389 longitudinal T1-weighted MRIs (at least two visits per subject) from ADNI, which consisted of 185 NC (age: 75.57 ± 5.06 years), 119 subjects with AD (age: 75.17 ± 7.57 years), 193 subjects diagnosed with sMCI (age: 75.63 ± 6.62 years), and 135 subjects diagnosed with pMCI (age: 75.91 ± 5.35 years). There was no significant age difference between the NC and AD cohorts ($p = 0.55$, two-sample t -test) as well as the sMCI and pMCI cohorts ($p = 0.75$). All longitudinal MRIs were preprocessed by a pipeline composed 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 to z-scores.

Implementation Details. Let C_k denote a Convolution(kernel size of $3 \times 3 \times 3$)-BatchNorm-LeakyReLU(slope of 0.2)-MaxPool(kernel size of 2) block with k filters, and CD_k an Convolution-BatchNorm-LeakyReLU-Upsample block. The architecture was designed as C_{16} - C_{32} - C_{64} - C_{16} - CD_{64} - CD_{32} - CD_{16} - CD_{16} with a convolution layer at the top for reconstruction. The regularization weights were set to $\lambda = 1.0$. The networks were trained for 50 epochs by the Adam optimizer with learning rate of 5×10^{-4} and weight decay of 10^{-5} . To make the algorithm computationally efficient, we built the graph dynamically on the mini-batch of each iteration. A batch size $N_{bs} = 64$ and neighbour size $N_{nb} = 5$ were used.

Evaluation. Five-fold cross-validation (folds split based on subjects) was conducted with 10% training subjects used for validation. Random flipping of brain hemispheres, and random rotation and shift were used as augmentation during training. We first qualitatively illustrated the trajectory vector field (Δz) in 2D space by projecting the 1024-dimensional bottleneck representations (z^t and z^s) to their first two principal components. We then estimated the global trajectory of the vector field by a curve fitted by robust linear mixed effect model, which

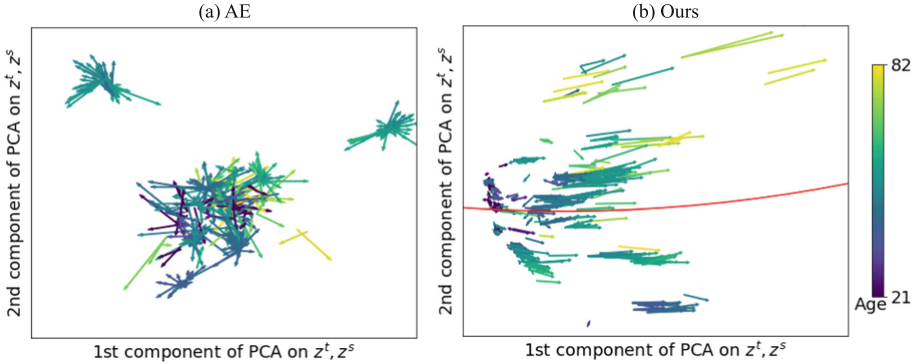


Fig. 2. Experiments on healthy aging: Latent space of AutoEncoder (AE) (a) and the proposed LNE (b) projected into 2D PCA space of z^t and z^s . Arrows represent Δz and are color-coded by the age of z^t . The global trajectory in (b) is fitted by robust linear mixed effect model (red curve). (Color figure online)

considered a quadratic fixed effect with random effect of intercepts. We further quantitatively evaluated the quality of the representations by using them for downstream tasks. Note, for these experiments, we removed the decoder and only kept the encoder with its pre-trained weights for the downstream tasks. On the dataset of healthy subjects, we used the representation z to predict the chronological age of each MRI to show that our latent space was stratified by age. Note, learning a prediction model for normal aging is an emerging approach for understanding structural changes of the human brain and quantifying impact of neurological diseases (e.g. estimating brain age gap [20]). R2 and root-mean-square error (RMSE) were used as accuracy metrics. For ADNI, we predicted the diagnosis group associated with each image pair based on both z and trajectory vector Δz to highlight the aging speed between visits (an important marker for AD). In addition to classifying NC and AD, we also aimed to distinguish pMCI from sMCI, a significantly more challenging classification task.

The classifier was designed as a multi-layer perceptron containing two fully connected layers of dimension 1024 and 64 with LeakyReLU activation. In a separate experiment, we fine-tuned the LNE representation by incorporating the encoder into the classification models. We compared the BACC (accounting for different number of training samples in each cohort) to models using the same architecture with encoders pre-trained by other representation learning methods, including unsupervised methods (AE, VAE [12]), self-supervised method (SimCLR [6]. Images of two visits of the same subject with simple shift and rotation augmentation were used as a positive pair in SimCLR), and longitudinal self-supervised method (LSSL [24]).

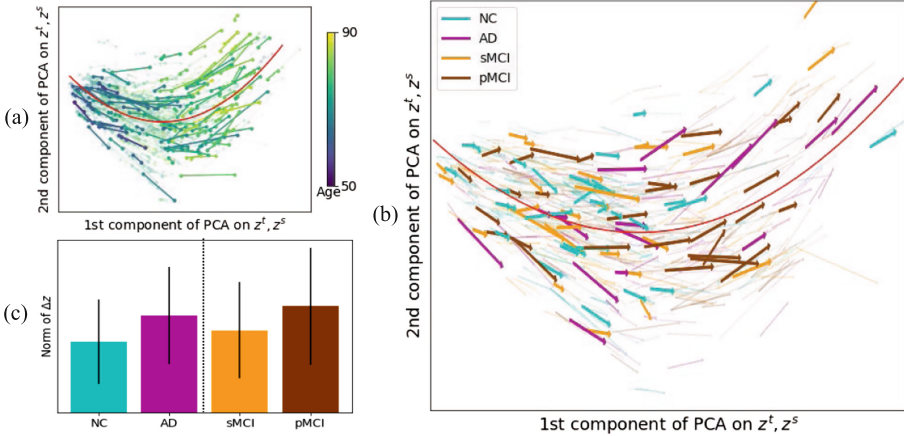


Fig. 3. Experiments on ADNI: (a) The age distribution of the latent space. Lines connecting z^t and z^s are color-coded by the age of z^t ; Red curve is the global trajectory fitted by a robust linear mixed effect model. (b) Trajectory vector field color-coded by diagnosis groups; (c) The norm of Δz encoding the speed of aging for 4 different diagnosis groups. (Color figure online)

3.1 Healthy Aging

Figure 2 illustrates the trajectory vector field derived on one of the 5 folds by the proposed method (Fig. 2(b)). We observe LNE resulted in a smooth vector field that was in line with the fitted global trajectory shown by the red curve in Fig. 2(b). Moreover, chronological age associated with the vectors (indicated by the color) gradually increased along the global trajectory (red curve), indicating the successful disentanglement of the aging effect in the latent space. Note, such continuous variation in the whole age range from 20 to 90 was solely learned by self-supervised training on image pairs with an average age interval of 3.8 years (without using their age information). Interestingly, the length of the vectors tended to increase along the global trajectory, suggesting a faster aging speed for older subjects. On the contrary, without regularizing the longitudinal changes, AE did not lead to clear disentanglement of brain age in the space (Fig. 2(a)).

As shown in Table 1 (left), we utilized the latent representation z to predict the chronological age of the subject. In the scenario that froze the encoder, the proposed method achieved the best performance with an R2 of 0.62, which was significantly better ($p < 0.01$, t-test on absolute errors) than the second-best method LSSL with an R2 of 0.59. In addition to R2, the RMSE metrics are given in the supplement Table S1, which also suggests that LNE achieved the most accurate prediction. These results align with the expectation that a pre-trained self-supervised model with explicitly modeling of aging effect can lead to better downstream supervised age prediction. Lastly, when we fine-tuned the encoder during training, LNE remained as the most accurate method (both LNE and LSSL achieved an R2 of 0.74).

Table 1. Supervised downstream tasks in frozen or fine-tune scenarios. Left: Age regression on healthy subjects with R2 as an evaluation metric. Right: classification on ADNI dataset with BACC as the metric.

Methods	Health aging (R2)		ADNI (BACC)			
	Age		NC vs AD		sMCI vs pMCI	
	Frozen	Fine-tune	Frozen	Fine-tune	Frozen	Fine-tune
No pretrain	–	0.72	–	79.4	–	69.3
AE	0.53	0.69	72.2	80.7	62.6	69.5
VAE [12]	0.51	0.69	66.7	77.0	61.3	63.8
SimCLR [6]	0.56	0.73	72.9	82.4	63.3	69.5
LSSL [24]	0.59	0.74	74.2	82.1	69.4	71.2
Ours (LNE)	0.62	0.74	81.9	83.6	70.6	73.4

3.2 Progression of Alzheimer’s Disease

We also evaluated the proposed method on the ADNI dataset. All 4 cohorts (NC, sMCI, pMCI, AD) were included in the training of LNE as the method was impartial to diagnosis groups (did not use labels for training). Similar to the results of the prior experiment, the age distribution in the latent space in Fig. 3(a) suggests a continuous variation with respect to brain development along the global trajectory shown by the red curve. We further illustrated the trajectory vector field by diagnosis groups in Fig. 3(b). While the starting points (z^t) of different diagnosis groups mixed uniformly in the field, vectors of AD (pink) and pMCI (brown) were longer than NC (cyan) and sMCI (orange). This suggests that LNE stratified the cohorts by their ‘speed of aging’ rather than age itself, highlighting the importance of using longitudinal data for analyzing AD and pMCI. This observation was also evident in Fig. 3(c), where AD and pMCI had statistically larger norm of Δz than the other two cohorts (both with $p < 0.01$). This finding aligned with previous AD studies [22] suggesting that AD group has accelerated aging effect compared to the NC group, and so does the pMCI group compared to the sMCI group.

The quantitative results on the downstream supervised classification tasks are shown in Table 1 (right). As the length of Δz was shown to be informative, we concatenated z^t with Δz as the feature for classification (classification accuracy based on z^t only is reported in the supplement Table S2). The representations learned by the proposed method yielded significantly more accurate predictions than all baselines ($p < 0.01$, DeLong’s test). Note that the accuracy of our model with the frozen encoder even closely matched up to other methods after fine-tuning. This was to be expected because only our method and LSSL explicitly modeled the longitudinal effects which led to more informative Δz . In addition, our method that focused on local smoothness could capture the potentially non-linear effects underlying the morphological change along time, while the ‘global linearity’ assumption in LSSL may lead to information loss in

the representations. It is worth mentioning that reliably distinguishing the subjects that will eventually develop AD (pMCI) from other MCI subjects (sMCI) is crucial for timely treatment. To this end, Supplement Table S3 suggests LNE improved over prior studies in classifying sMCI vs. pMCI, highlighting potential clinical values of our method. Ablation study on two important hyperparameters N_{nb} and λ is reported in the supplement Table S4.

4 Conclusion

In this work, we proposed a self-supervised representation learning framework, called LNE, that incorporates advantages from the repeated measures design in longitudinal neuroimaging studies. By building the dynamic graph and learning longitudinal neighbourhood embedding, LNE yielded a smooth trajectory vector field in the latent space, while maintaining a globally consistent progression trajectory that modeled the morphological change of the cohort. It successfully modeled the aging effect on healthy subjects, and enabled better chronological age prediction compared to other self-supervised methods. Although LNE was trained without the use of diagnosis labels, it demonstrated capability of differentiating diagnosis groups on the ADNI dataset based on the informative trajectory vector field. When evaluated for downstream task of classification, it showed superior quantitative classification performance as well.

Acknowledgement. This work was supported by NIH funding R01 MH113406, AA017347, AA010723, and AA005965.

References

1. Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S.: Social LSTM: human trajectory prediction in crowded spaces. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 961–971 (2016)
2. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: An unsupervised learning model for deformable medical image registration. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 9252–9260 (2018)
3. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE Trans. Med. Imaging* **38**(8), 1788–1800 (2019)
4. Carass, A., et al.: Longitudinal multiple sclerosis lesion segmentation: resource and challenge. *NeuroImage* **148**, 77–102 (2017)
5. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882* (2020)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International Conference on Machine Learning, pp. 1597–1607. PMLR (2020)

7. Dalca, A.V., Yu, E., Golland, P., Fischl, B., Sabuncu, M.R., Eugenio Iglesias, J.: Unsupervised deep learning for Bayesian brain MRI segmentation. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11766, pp. 356–365. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32248-9_40
8. Garcia, T.P., Marder, K.: Statistical approaches to longitudinal data analysis in neurodegenerative diseases: Huntington’s disease as a model. *Curr. Neurol. Neurosci. Rep.* **17**(2), 14 (2017)
9. Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S., Alahi, A.: Social GAN: Socially acceptable trajectories with generative adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2255–2264 (2018)
10. Hassani, K., Khasahmadi, A.H.: Contrastive multi-view representation learning on graphs. In: International Conference on Machine Learning, pp. 4116–4126. PMLR (2020)
11. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9729–9738 (2020)
12. Kingma, D.P., Welling, M.: Auto-encoding variational Bayes. arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (2013)
13. Klicpera, J., Weissenberger, S., Günnemann, S.: Diffusion improves graph learning. arXiv preprint [arXiv:1911.05485](https://arxiv.org/abs/1911.05485) (2019)
14. Kolesnikov, A., Zhai, X., Beyer, L.: Revisiting self-supervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1920–1929 (2019)
15. Li, H., Fan, Y.: Non-rigid image registration using self-supervised fully convolutional networks without training data. In: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018), pp. 1075–1078. IEEE (2018)
16. Liu, X., Zhang, F., Hou, Z., Wang, Z., Mian, L., Zhang, J., Tang, J.: Self-supervised learning: Generative or contrastive **1**(2). arXiv preprint [arXiv:2006.08218](https://arxiv.org/abs/2006.08218) (2020)
17. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint [arXiv:1807.03748](https://arxiv.org/abs/1807.03748) (2018)
18. Ouyang, J., et al.: Longitudinal pooling & consistency regularization to model disease progression from MRIs. *IEEE J. Biomed. Health Inform.* (2020)
19. Sabokrou, M., Khalooei, M., Adeli, E.: Self-supervised representation learning via neighborhood-relational encoding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 8010–8019 (2019)
20. Smith, S.M., et al.: Brain aging comprises many modes of structural and functional change with distinct genetic and biophysical associations. *Elife* **9**, e52677 (2020)
21. Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. arXiv preprint [arXiv:1906.05849](https://arxiv.org/abs/1906.05849) (2019)
22. Toepper, M.: Dissociating normal aging from alzheimer’s disease: a view from cognitive neuroscience. *J. Alzheimer’s Dis.* **57**(2), 331–352 (2017)
23. Whitwell, J.L.: Longitudinal imaging: change and causality. *Curr. Opin. Neurol.* **21**(4), 410–416 (2008)
24. Zhao, Q., Liu, Z., Adeli, E., Pohl, K.M.: LSSL: Longitudinal self-supervised learning. arXiv preprint [arXiv:2006.06930](https://arxiv.org/abs/2006.06930) (2020)