



An Image Feature Mapping Model for Continuous Longitudinal Data Completion and Generation of Synthetic Patient Trajectories

Clément Chadebec¹, Evi M. C. Huijben²(✉), Josien P. W. Pluim²,
Stéphanie Allassonnière¹, and Maureen A. J. M. van Eijnatten²

¹ Université Paris Cité, INRIA, Inserm, Sorbonne Université, Paris, France
clement.chadebec@inria.fr

² Department of Biomedical Engineering, Medical Image Analysis Group,
Eindhoven University of Technology, Eindhoven, The Netherlands
e.m.c.huijben@tue.nl

Abstract. Longitudinal medical image data are becoming increasingly important for monitoring patient progression. However, such datasets are often small, incomplete, or have inconsistencies between observations. Thus, we propose a generative model that not only produces continuous trajectories of fully synthetic patient images, but also imputes missing data in existing trajectories, by estimating realistic progression over time. Our generative model is trained directly on features extracted from images and maps these into a linear trajectory in a Euclidean space defined with velocity, delay, and spatial parameters that are learned directly from the data. We evaluated our method on toy data and face images, both showing simulated trajectories mimicking progression in longitudinal data. Furthermore, we applied the proposed model on a complex neuroimaging database extracted from ADNI. All datasets show that the model is able to learn overall (disease) progression over time.

Keywords: Longitudinal data · Generative model · Synthetic images

1 Introduction

Longitudinal medical image data are important for *e.g.* modelling disease progression [1, 27] or monitoring treatment response [3]. However, such datasets often suffer from incomplete or inconsistent observations, and are often limited in terms of size, diversity, and balance. Generally, using inadequate data can lead to

C. Chadebec and E. M. C. Huijben—Equal contribution.

S. Allassonnière and M. A. J. M. van Eijnatten—Equal contribution.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-18576-2_6.

poor performances when being used to train machine learning (ML) models [23] for medical image analysis tasks such as classification [26] or segmentation [15].

To increase the size and variability of (non-longitudinal) medical imaging datasets, conventional data augmentation techniques such as rotation, cropping, or more resourceful augmentations [11] have been widely used [25]. However, the improved performances of deep generative models have given them the potential to perform image synthesis. Examples of such models are Generative Adversarial Networks (GANs) [10], which generate realistic images using a discriminator that distinguishes between real and synthetic images, and Variational Autoencoders (VAEs) [14], which constrain image features to follow a given prior distribution in order to generate synthetic images. These models have shown potential for synthesizing medical images of various modalities such as magnetic resonance imaging (MRI) [5, 6, 24], computed tomography (CT) [8, 22], X-ray [18, 21], or positron emission tomography (PET) [2]. In addition, several methods have been proposed to address data imputation or progression modelling in longitudinal imaging data of *e.g.* MRI [13, 17] or simulated discrete progressions [20].

Although the topics of inter- and extrapolating longitudinal (medical) imaging data are well studied, to the best of our knowledge there is no model that addresses both of these aspects at once and is able to continuously generate realistic trajectories. In this paper, we propose a new deep generative model that is capable of: (1) generating realistic progression in images, (2) imputing missing data in existing patient trajectories, and (3) producing synthetic images with corresponding trajectories of non-existent patients¹.

2 Proposed Method

We propose a new generative model for longitudinal imaging data that consists of two steps. In the first step, relevant features are extracted from the input images using a VAE, and the second step maps these features into a linear trajectory to account for the progression over time. In the following, we refer to an observation, *e.g.* an image, as $y_{i,j} \in \mathcal{Y}$, with $i \in [1, N]$ the individual’s identifier, $t_j \in \mathbb{R}_+^*$, where $j \in [0, P_i]$ the time of the observation. N is the number of individuals and P_i is the number of observations of i after the first time visit t_0 .

2.1 Feature Extraction

Medical images are often complex and high-dimensional data. Therefore, instead of proposing a model directly acting on images, we propose to first extract meaningful features using a VAE (referred to as *the VAE* in the following). We use an autoencoder because it constrains comparable images to be encoded into similar locations such that minor variations in the latent space lead to smooth transformations in the image space. Since we expect smooth progressions, the VAE is likely to directly unveil trajectories in the latent space, thereby facilitating the second step of our method (referred to as *the generative model* in the following). In the following $x_{i,j} \in \mathcal{M}$ refers to the features of observation $y_{i,j}$.

¹ Code and dataset details are available at <https://github.com/evihuijben/longVAE>.




Fig. 1. Model sketch. First, features are extracted from images using the VAE (step 1), then, the proposed generative model maps these features to a straight line in Euclidean space (step 2). Network details are provided in Appendix 2.

2.2 Trajectory Modelling

We propose to learn parametric functions that map the features onto a linear trajectory in a d -dimensional Euclidean space \mathbb{R}^d with standard basis $\{\mathbf{e}_1, \dots, \mathbf{e}_d\}$, accounting for an individual’s progression. We use the framework proposed in [17], in which an individual’s progression trajectory at time t is modelled in \mathbb{R}^d as

$$l_i(t) = \exp(\eta_i)(t - \tau_i) \cdot \mathbf{e}_1 + \sum_{k=2}^d \lambda_i^k \cdot \mathbf{e}_k, \quad (1)$$

where η_i is a velocity parameter, τ_i is a delay, and $\lambda_i = (\lambda_i^k)_{2 \leq k \leq d}$ are spatial parameters. Contrary to [17], we adopt a fully variational approach to make the model generative in a similar fashion as [14]. Assuming a set of embeddings $x = \{(x_{i,j})_{1 \leq i \leq N, 0 \leq j \leq P_i}\} \in \mathcal{M}$, we first assume that given two individuals i and i' , the features $x_{i,j}$ and $x_{i',j}$ are independent. Therefore, we propose to maximise the following likelihood objective $p(x) = \prod_{i=1}^N p(x_i)$, where $x_i = (x_{i,0}, \dots, x_{i,P_i})$. We further assume that the latent variables $z_i = (\eta_i, \tau_i, \lambda_i^2, \dots, \lambda_i^d) \in \mathbb{R}^{d+1}$ in Eq. (1) are such that the features of individual i at time t_j are generated by:

$$p_{\theta}(x_{i,j}|z_i) = \mathcal{N}(\mu_{\theta}(l_i(t_j)), \sigma \cdot I_d), \quad (2)$$

where $l_i(t_j)$ is the linear trajectory evaluated at t_j , and $\mu_{\theta} : \mathbb{R}^d \rightarrow \mathcal{M}$ is parameterised using a multilayer perceptron (MLP) and maps \mathbb{R}^d to the feature space. The variation introduced by the stochastic model is the d -dimensional unit matrix I_d multiplied by a positive constant σ . We further assume that η_i , τ_i , and λ_i are independent and that for a given individual i , the features $x_{i,j}$ taken conditionally to z_i are independent. Furthermore, prior distributions over the latent variables are: $\eta_i \sim \mathcal{N}(0, \sigma_{\eta})$, $\tau_i \sim \mathcal{N}(0, \sigma_{\tau})$, $\lambda_i \sim \mathcal{N}(0, I_{d-1})$, with the dataset dependent priors σ_{η} and σ_{τ} . Finally, the likelihood for an individual i is:

$$p(x_i) = \int_{z_i \in \mathbb{R}^{d+1}} p_\theta(x_i|z_i)p(z_i)dz_i = \int_{z_i \in \mathbb{R}^{d+1}} \prod_{j=0}^{P_i} p_\theta(x_{i,j}|z_i) \prod_{\kappa_i \in \{\eta_i, \tau_i, \lambda_i\}} p(\kappa_i)d\kappa_i. \tag{3}$$

Since $p(z_i|x_i)$, the true posterior distribution, is unknown, we rely on variational inference [12]. Hence, we introduce a variational distribution $q_\varphi(z_i|x_i) = q_\varphi(\eta_i|x_i)q_\varphi(\tau_i|x_i)q_\varphi(\lambda_i|x_i)$ and derive a new estimate of the likelihood $p(x_i) = \mathbb{E}_{z_i \sim q_\varphi(z_i|x_i)} \left[\frac{p(x_i, z_i)}{q_\varphi(z_i|x_i)} \right]$. We then compute a lower bound on the true objective using Jensen inequality and importance sampling using the variational distribution.

$$\log p(x_i) \geq \mathbb{E}_{z_i \sim q_\varphi(z_i|x_i)} \left[\log p(x_i|z_i) \right] - \sum_{\kappa_i \in \{\eta_i, \tau_i, \lambda_i\}} KL(q_\varphi(\kappa_i|x_i)|p(\kappa_i)), \tag{4}$$

with KL the Kullback-Leibler divergence. In practice, we use multivariate Gaussians as variational distributions: $\eta_i \sim \mathcal{N}(\mu_\varphi^{\eta_i}, \Sigma_\varphi^{\eta_i})$, $\tau_i \sim \mathcal{N}(\mu_\varphi^{\tau_i}, \Sigma_\varphi^{\tau_i})$ and $\lambda_i \sim \mathcal{N}(\mu_\varphi^{\lambda_i}, \Sigma_\varphi^{\lambda_i})$. The parameters for progression, η_i and τ_i , are estimated from an input sequence using a recurrent neural network (RNN), while the spatial parameters, $(\lambda_i^2, \dots, \lambda_i^d)$, are computed from the features of the image acquired at time t_0 , which are estimated by the first MLP. The implementation details of the RNN and MLP can be found in Appendix 2, and a sketch of the model is presented in Fig. 1. Taking only the first image’s features for the spatial parameters allows to estimate their value even if only one observation is available and to generate possible future progressions. Finally, we obtain the following loss function for one individual (removing constant terms):

$$\mathcal{L}_i = \sum_{j=0}^{P_i} \|x_{i,j} - \mu_\theta(l_i(t_j))\|^2 + \sum_{\kappa_i \in \{\eta_i, \tau_i, \lambda_i\}} KL(q_\varphi(\kappa_i|x_i)|p(\kappa_i)). \tag{5}$$

After training, we can either 1) generate fully synthetic trajectories using the aforementioned prior distributions, 2) produce possible progressions for a given individual i by estimating its λ_i and varying η_i and τ_i , or 3) interpolate and extrapolate existing trajectories by estimating the latent variables. Image sequences are then generated by recovering the features corresponding to a linear trajectory evaluated at a given time using a second MLP and passing them to the decoder of the VAE. In practice, we sample λ_i with a mixture of Gaussians since [9] recently showed that this approach alleviates the low expressiveness of the prior and allows to generate more convincing samples.

3 Data

We evaluate the proposed model using three longitudinal datasets. The first dataset is a toy dataset referred to as Starmen² [7] consisting of 64×64 binary

² Downloaded from <https://doi.org/10.5281/zenodo.5081988>.

images of 1,000 individuals that portray synthetic transformations based on the longitudinal model of [4], captured in 10 observations per individual. The second dataset, CelebA (aligned and cropped version downloaded in 2021) [16], consists of 64×64 RGB images of celebrities’ faces. To resemble longitudinal medical images, we converted these images to grayscale and applied a simulated progression model by applying a non-linear intensity transform, a growth factor, a rotation, and adding Gaussian noise. This dataset can be considered very challenging since the images undergo global and local geometric transformations and photometric variations. The last dataset was obtained from the Alzheimer’s Disease Neuroimaging Initiative³. We used a total of 8,318 MRI scans, obtained from 1,799 subjects, with an average of 4.6 ± 2.3 scans per person. The average time between the first and the last scan was 2.9 ± 2.4 years. We selected the 100th axial slice of every preprocessed scan and cropped it to 182×182 . The subject’s ages were used to define the observation times for the generative model, which were normalised between the overall oldest and youngest age. Details of the datasets (*e.g.* preprocessing steps, progression model, data splits, and example image trajectories) can be found in Appendix 1.

4 Experiments

Most experiments in this section are performed using Starmen and CelebA because these datasets are fully controlled and allow visual evaluation by non-medical experts. ADNI is used to show that results can be extended to medical data. In what follows, the models are selected on the validation set and tested on an hold-out test set. Experimental and implementation details are provided in Appendix 2.

Feature Extraction and Reconstruction. First, we train the VAE on each training set, disregarding the longitudinal component, and confirm the hypothesis that the features directly unveil clear trajectories over time, as can be seen in Fig. 1b in Appendix 1. To justify that mapping those trajectories to linear ones (step 2 in Fig. 1) is not too constraining, we analyse the reconstruction results obtained by 1) only encoding and decoding test images using the VAE (*base*), and 2) training the generative model to map the extracted feature trajectories to straight lines (Eq. (1)), evaluate $l(t)$ at observation times and pass the corresponding features to the decoder of the VAE (*ours*). Figure 2a and 2c show the mean squared error (MSE) and structural similarity (SSIM), respectively, of the test set reconstructions. Note that the results obtained using the proposed model is not expected to be better than the one obtained using the

³ Data used in preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (<http://adni.loni.usc.edu>). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf.




Fig. 2. Mean and standard deviation of MSE/SSIM (a, b/c, d) for various evaluations. (a/c) Metric between consecutive images in the test sequences (*ref.*) and reconstruction metrics using only the VAE (*base*) or the generative model (*ours*). (b/d) Metrics for the next and last image extrapolated based on a varying input sequence length.

VAE (*base*) because the generative model only acts on the features and we do not use any image-based reconstruction cost during its training. The metric values can be put into perspective by considering the mean value between two consecutive images in the test set (*ref.*). The visual reconstructions in the second row of Fig. 3 show that linear trajectory modelling does not considerably affect the image reconstruction ability of the model.




Fig. 3. Extrapolation of different test input sequences for Starmen (left) and CelebA (right). The first two rows represent the ground truth and reconstructions (*ours*), respectively. Red squares highlight images that were not provided to the model. Deviation from the true test Starmen image is presented in colour. (Color figure online)

Trajectory Extrapolation. In this section, we investigate whether the proposed model is able to extrapolate realistic trajectories from existing input data. To do so, we use the same model as before, but only provide the model with an image sequence of varying length and assess its ability to reconstruct either the next or the last image in the sequence. Figure 2b and Fig. 2d show the MSE and SSIM, respectively, of the ground truth and the extrapolated images based on a varying input sequence length. It can be seen that extrapolations become more reliable when a longer input sequence is given. This can also be observed from the visuals in Fig. 3, which show larger deviations from the ground truth when fewer images are presented. This experiment shows that in each case the model is able to estimate the progression: the left arm of the Starmen is raising and the CelebA head rotates, becomes bigger and contrast changes as expected. However, the model seems to underestimate the trajectory velocity as the input sequence

becomes shorter. This aspect could potentially be mitigated by training using sequences of different lengths.

Data Imputation. We validate the ability of the model to impute missing data using input sequences simulating partial patient follow-ups. We simulate this by removing 50% of the training, validation, and test data acquired after t_0 using the Starmen and CelebA datasets. The VAE is trained using the 50% available images, after which the generative model learns to map the features onto a linear trajectory. In Fig. 4 we show the reconstructed samples at observation times.




Fig. 4. Data imputation in test sequences with 50% missing data after t_0 . Top rows show ground truth trajectories, red squares represent imputed images. (Color figure online)

Trajectory Generation. We also demonstrate that the proposed model can generate synthetic trajectories. We consider two cases: generating possible trajectories for a single image acquired at t_0 and generating a fully synthetic trajectory based on a synthetic image at t_0 . In the first case, we first recover λ_i by encoding the real image using the VAE, estimate its value using the generative model and then sample η and τ from their priors as described in Sect. 2.2 and Appendix 2. In the second case, we first generate a synthetic λ and sample η and τ as aforementioned. To demonstrate the differences in these parameters, Fig. 5 shows trajectories obtained with varying delay τ (a) and velocity η (b), possible trajectories from an input image (c) and fully synthetic trajectories (d).

Real images are extracted from the test set and highlighted with blue frames. The results show that the proposed model allows to decorrelate spatial (λ) and time parameters (η and τ) since all images in a trajectory represent the same individual that undergoes smooth progressive change.

Neuroimaging Data. Finally, we validate the ability of the model to generate Alzheimer’s disease progression trajectories. Figure 5e and 5f show trajectories generated from an existing input image and a synthetic image, respectively. The generated trajectories appear realistic because the ventricles grow over time, which is a marker of ageing and Alzheimer’s disease progression [19]. Moreover, the proposed model seems to preserve the morphology represented at the first time point for both real and fake subjects. However, the generated disease progression trajectories still need to be assessed in more detail, for example by means of visual analysis by a medical expert or by training a deep learning-based classifier. Beside generating synthetic trajectories, we also investigate the extrapolation capability of the proposed model for the ADNI data, which is shown in

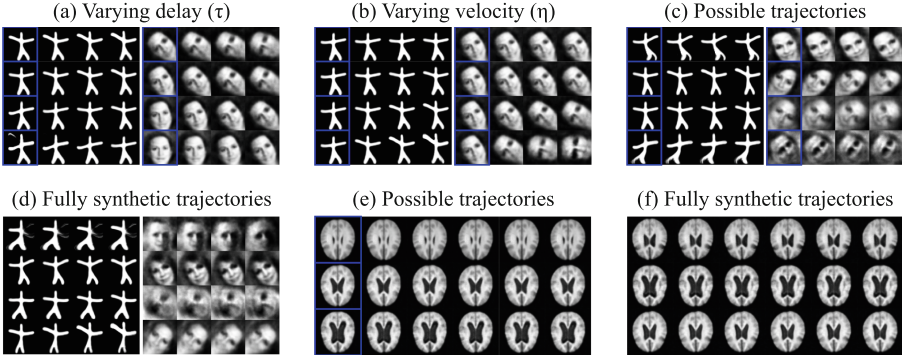


Fig. 5. Synthetic trajectories derived from real images (indicated by blue frames): (a-c, e) or synthetic images (d, f). (Color figure online)

Fig. 6. Contrary to the Starmen and CelebA experiments, this experiment shows a better performance for a shorter input sequence length. Generally, the proposed model seems to underestimate the disease progression (as estimated by η and τ), leading to a worse quantitative result for a later extrapolated sample.

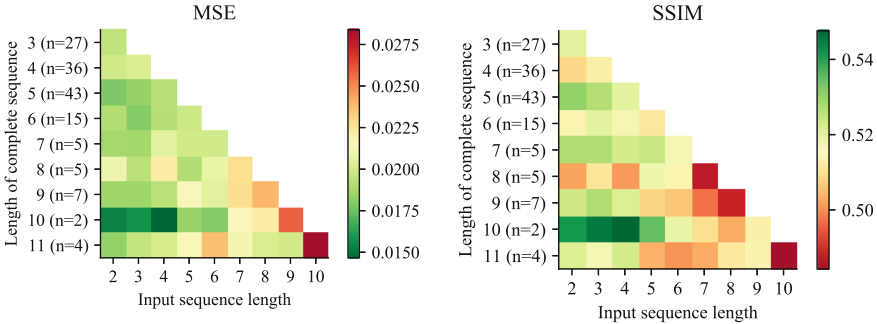


Fig. 6. Mean MSE (left) and SSIM (right) for the extrapolated next image after a given input sequence of the ADNI test set, with n the number of subjects. For interpretation of the colour bars, the reader is referred to the online version. (Color figure online)

5 Discussion and Conclusion

In this study we proposed a new continuous generative model capable of synthesising longitudinal imaging data to perform trajectory extrapolation, data imputation and smooth and probable synthetic trajectory generation. A notable strength of our model lies in its *two-step* architecture, which allows substituting the VAE to make the model suitable for any data type, *e.g.* using clinical scores directly as features. We believe that this work is a step towards synthesis and augmentation of longitudinal medical (image) datasets. However, the

model needs more optimisation for such a high-dimensional complex medical imaging dataset, and a better trade-off between dimensionality reduction and efficient training of the generative model should be investigated. Furthermore, the hypothesis of smooth trajectories could be put into perspective by considering the disentangled ‘brain age’ instead of the real patient’s age [27]. Future work should also focus on validating the ability of the model to perform reliable data augmentation for ML-based classification tasks or assess its relevance to perform treatment response analysis.

References

1. Aghili, M., Tabarestani, S., Adjouadi, M., Adeli, E.: Predictive modeling of longitudinal data for Alzheimer’s disease diagnosis using RNNs. In: Rekek, I., Unal, G., Adeli, E., Park, S.H. (eds.) PRIME 2018. LNCS, vol. 11121, pp. 112–119. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00320-3_14
2. Bi, L., Kim, J., Kumar, A., Feng, D., Fulham, M.: Synthesis of positron emission tomography (PET) images via multi-channel generative adversarial networks (GANs). In: Cardoso, M.J., et al. (eds.) CMMI/SWITCH/RAMBO 2017. LNCS, vol. 10555, pp. 43–51. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67564-0_5
3. Blackledge, M.D., et al.: Assessment of treatment response by total tumor volume and global apparent diffusion coefficient using diffusion-weighted MRI in patients with metastatic bone disease: a feasibility study. PLoS ONE **9**(4), e91779 (2014)
4. Bône, A., Colliot, O., Durrleman, S.: Learning distributions of shape trajectories from longitudinal datasets: a hierarchical model on a manifold of diffeomorphisms. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 9271–9280 (2018)
5. Calimeri, F., Marzullo, A., Stamile, C., Terracina, G.: Biomedical data augmentation using generative adversarial neural networks. In: Lintas, A., Rovetta, S., Verschure, P.F.M.J., Villa, A.E.P. (eds.) ICANN 2017. LNCS, vol. 10614, pp. 626–634. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68612-7_71
6. Chadebec, C., Thibeau-Sutre, E., Burgos, N., Allasonnière, S.: Data augmentation in high dimensional low sample size setting using a geometry-based variational autoencoder. IEEE Trans. Pattern Anal. Mach. Intell. (2022)
7. Couronné, R., Vernhet, P., Durrleman, S.: Longitudinal self-supervision to disentangle inter-patient variability from disease progression. In: de Bruijne, M., et al. (eds.) MICCAI 2021. LNCS, vol. 12902, pp. 231–241. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87196-3_22
8. Frid-Adar, M., Diamant, I., Klang, E., Amitai, M., Goldberger, J., Greenspan, H.: GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification. Neurocomputing **321**, 321–331 (2018)
9. Ghosh, P., Sajjadi, M.S., Vergari, A., Black, M., Schölkopf, B.: From variational to deterministic autoencoders. In: International Conference on Learning Representations (ICLR) (2020)
10. Goodfellow, I., et al.: Generative adversarial nets. In: Advances in Neural Information Processing Systems, vol. 27 (2014)
11. Hussain, Z., Gimenez, F., Yi, D., Rubin, D.: Differential data augmentation techniques for medical imaging classification tasks. In: AMIA Annual Symposium Proceedings, vol. 2017, p. 979. American Medical Informatics Association (2017)

12. Jordan, M.I., Ghahramani, Z., Jaakkola, T.S., Saul, L.K.: An introduction to variational methods for graphical models. In: Jordan, M.I. (ed.) *Machine Learning*. NATO ASI Series, pp. 105–161. Springer, Dordrecht (1998). https://doi.org/10.1007/978-94-011-5014-9_5
13. Kim, S.T., Küçükaşlan, U., Navab, N.: Longitudinal brain MR image modeling using personalized memory for Alzheimer’s disease. *IEEE Access* **9**, 143212–143221 (2021)
14. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (2013)
15. Liu, X., Song, L., Liu, S., Zhang, Y.: A review of deep-learning-based medical image segmentation methods. *Sustainability* **13**(3), 1224 (2021)
16. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (2015)
17. Louis, M., Couronné, R., Koval, I., Charlier, B., Durrleman, S.: Riemannian geometry learning for disease progression modelling. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) *IPMI 2019*. LNCS, vol. 11492, pp. 542–553. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20351-1_42
18. Madani, A., Moradi, M., Karargyris, A., Syeda-Mahmood, T.: Chest X-ray generation and data augmentation for cardiovascular abnormality classification. In: *Medical Imaging 2018: Image Processing*, vol. 10574, pp. 415–420. International Society for Optics and Photonics, SPIE (2018)
19. Nestor, S.M., et al.: Ventricular enlargement as a possible measure of Alzheimer’s disease progression validated using the Alzheimer’s disease neuroimaging initiative database. *Brain* **131**(9), 2443–2454 (2008)
20. Ramchandran, S., Tikhonov, G., Kujanpää, K., Koskinen, M., Lähdesmäki, H.: Longitudinal variational autoencoder. In: *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*. *Proceedings of Machine Learning Research*, vol. 130, pp. 3898–3906. PMLR (2021)
21. Salehinejad, H., Valaee, S., Dowdell, T., Colak, E., Barfett, J.: Generalization of deep neural networks for chest pathology classification in X-rays using generative adversarial networks. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 990–994 (2018)
22. Sandfort, V., Yan, K., Pickhardt, P.J., Summers, R.M.: Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks. *Sci. Rep.* **9**(1), 16884 (2019)
23. Shin, H.C., et al.: Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. *IEEE Trans. Med. Imaging* **35**(5), 1285–1298 (2016)
24. Shin, H.-C., et al.: Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In: Gooya, A., Goksel, O., Oguz, I., Burgos, N. (eds.) *SASHIMI 2018*. LNCS, vol. 11037, pp. 1–11. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00536-8_1
25. Shorten, C., Khoshgoftaar, T.M.: A survey on image data augmentation for deep learning. *J. Big Data* **6**(1), 60 (2019)
26. Wen, J., et al.: Convolutional neural networks for classification of Alzheimer’s disease: overview and reproducible evaluation. *Med. Image Anal.* **63**, 101694 (2020)
27. Zhao, Q., Liu, Z., Adeli, E., Pohl, K.M.: Longitudinal self-supervised learning. *Med. Image Anal.* **71**, 102051 (2021)