Longitudinal Variational Autoencoders learn a Riemannian progression model for imaging data.

Benoît Sauty
Stanley Durrleman

Abstract

Interpretable progression models for longitudinal neuroimaging data are crucial to understanding neurodegenerative diseases. Well validated geometric progression models for biomarkers do not scale for such high dimensional data. In this work, we analyse a recent approach that combines a Variational Autoencoder with a latent linear mixed-effects model, and demonstrate that imposing a Euclidean prior on the latent space allows the network to learn the geometry of the observation manifold, and model non linear dynamics.
Keywords: Longitudinal progression - Riemannian manifolds - Variational Autoencoders

1. Introduction

Many flavours of geometric disease progression models for biomarkers assume that the data lie on a manifold and evolve according to non-linear geodesics across time. Generalization of such models to high dimensional neuroimaging data has called for the use of deep learning dimension reduction techniques. On the other hand, much attention has been put on finding geometric interpretations of deep learning generative models, in order to refine the learned distribution. Through the geometric interpretation of a recently introduced Longitudinal Variational Autoencoder model, we will bridge the gap between the well validated disease progression models for biomarkers and the novel deep-learning-based methods.

1.1. Geometric disease progression models

Mixed-effects models provide one of the most popular disease progression framework for longitudinal data. Individual trajectories are modeled as small variations around the population-average trajectory. Early models used linear modeling while non linearities were later added with polynomial, logistic and exponential regressions. One less restrictive assumption is to consider that the observed biomarkers follow continuous trajectories in the space of observations that is assumed to be a Riemannian manifold (Schiratti et al., 2015). This approach provides spatio-temporal models that describe the average trajectory as a geodesic and the inter-patients variability as the parallel transport of this trajectory on the manifold. Particular cases have been derived with a metric that is set *a priori* in order to yield closed-form trajectories (Koval et al., 2017; Schiratti et al., 2017), but less restrictive assumptions have been made in order to learn the appropriate metric from the data (Gruffaz et al., 2021; Sauty and Durrleman, 2022b), at the expense of an added computational burden for the computation of the exponentials as Hamiltonian flows.

Those models however do not scale with high dimensional data and are mostly limited to the study of biomarkers. For neurodegenerative diseases, studying the alterations of the brain is a crucial task in order to understand the early stages of the disease. However, progression models for neuroimaging data are still understudied, and the use of deep learning as a "black-box" constitutes a barrier to widespread application by clinicians.

1.2. Geometry-aware deep generative models

Deep generative models aim to provide a neural network generator $p_\theta : z \in \mathcal{Z} \mapsto x \in \mathcal{X}$, parametrized by θ, from a latent space $\mathcal{Z} \subset \mathbb{R}^n$ to a target space $\mathcal{X} \subset \mathbb{R}^N$ (usually $n << N$) such that p_θ maps a simple distribution in $\mathcal{Z}$ to the complex data distribution in $\mathcal{X}$. The Jacobian of the generator $\mathrm{J} = \frac{\partial p_\theta}{\partial z}$ provides a linear mapping from tangent vectors of $\mathcal{Z}$ to tangent vectors of $\mathcal{X}$ and $\mathrm{M} = \mathrm{J}^T\mathrm{J}$ thus defines a smoothly varying inner product on the tangent bundle $T\mathcal{Z}$, and can be seen as a Riemannian metric on $\mathcal{Z}$. For VAEs, the generator network – the decoder – is coupled with an encoder network q_ϕ, parametrized by ϕ, that provides a variational estimate of the posterior distribution $q_\phi(z|x)$ for $x \in \mathcal{X}$. When adding a prior on the latent space, we obtain a lower bound for the log likelihood that can be optimized through backpropagation (Kingma and Welling, 2013).

In Arvanitidis et al. (2017); Chen et al. (2018), the Riemannian metric M is used to describe the geometry of the latent space of a VAE, significantly improving interpolants and distance computation, and Shao et al. (2018) experimentally note that the curvature of the learned latent space is almost flat. Falorsi et al. (2018), on the other hand, proposed to introduce Riemannian geometry tools in the input space of the VAE itself, in order to allow the generation of manifold valued data.

1.3. Longitudinal Variational Autoencoders

In recent work, Sauty and Durrleman (2022a) introduced a progression model for imaging data that embeds a linear mixed-effect model in the latent space of a VAE. This model can then be calibrated on longitudinal databases of neuroimaging data such as t1-MRI or FDG-PET scans in order to provide a simple parametrization of the individual trajectories in $\mathcal{Z}$, that translate into non-linear trajectories in the image domain. An iterative MCMC optimization scheme allows to ensure that the VAE reconstructs images correctly while allowing to match the simple linear mixed-effect description of latent encodings to the complex trajectories in the observation space. Such an approach allows sampling individual trajectories at any timepoint – future or past – for prediction and missing data imputation.

Contributions The geometric nature of the VAE has been recognised and extensively studied. The most widespread paradigm is to embed the latent space with a Riemannian manifold structure, equipped with the pull-back of the Euclidean metric of the observation space trough the decoder. On the other hand, it is standard to consider that medical data lie on a non-Euclidean manifold, on which trajectories are parametrized using mixed-effects models, to provide interpretable non-linear progression models. For high dimensional neuroimaging data, the use of dimension reduction techniques such as VAEs is required.

In this work we provide a geometric interpretation of the LVAE progression model. We demonstrate that imposing a Euclidean prior on the latent space allows to characterise the Riemannian metric of the observation manifold as the push-forward of the Euclidean metric through the decoder, providing a simple way to compute exponentials and geodesics in the image domain. This "reversed" conception of the geometry of the VAEs, first introduced in Louis et al. (2019), bridges the gap between well-validated geometric progression models for biomakers and recent approaches that integrate deep learning tools in order to model longitudinal neuroimaging data.

2. Geometry of the observation manifold

With the former notations, $p_\theta(\mathcal{Z})$ is a n-dimensional immersed submanifold of $\mathbb{R}^N$ if the activations functions are smooth and monotonic and if the weight matrix of each layer has maximal rank (Shao et al., 2018). The first condition is a design choice and the second condition can be checked after training. If we notate g the Euclidean metric on $\mathcal{Z}$, the push-forward of g on $p_\theta(\mathcal{Z})$ is defined, for any smooth vector fields U, V on $p_\theta(\mathcal{Z})$, as

$$p_\theta^*(g)(U,V) = g((p_\theta)_*(U),(p_\theta)_*(V))$$

where $(p_\theta)_*(U)$ and $(p_\theta)_*(V)$ are the pull-back of U and V on $\mathcal{Z}$, which are defined by $(p_\theta)_*(U) : f \mapsto U(f \circ p_\theta^{-1})$ for smooth functions $f : \mathcal{Z} \mapsto \mathbb{R}$. Any geodesic $\gamma : [0,1] \to \mathcal{Z}$ on $(\mathcal{Z}, g)$ thus translates to a geodesic $p_\theta : [0,1] \to \mathcal{X}$ on $(p_\theta(\mathcal{Z}), p_\theta^*(g))$. The decoder network is an isometry between $(\mathcal{Z}, g)$ and $(p_\theta(\mathcal{Z}), p_\theta^*(g))$ so computation of Riemannian exponential or parallel transport can be done in $\mathcal{Z}$ inexpensively, before being push-forward to the observation space. In the context of LVAE, the linear trajectories of the latent space and inter-patients translations thus translate to geodesics and parallel transport in the corresponding submanifold of the observation space.

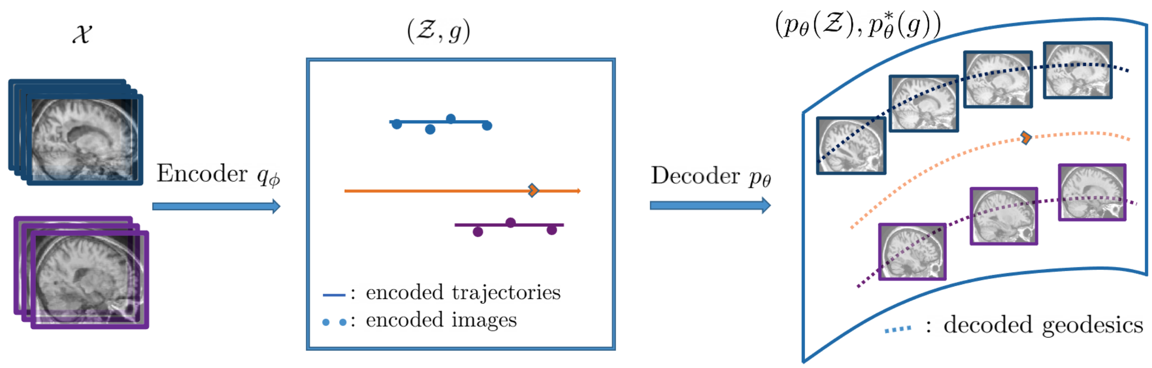

Figure 1: Schematic representation of the LVAE model. A linear trajectory is fitted on the latent representations (dots in $\mathcal{Z}$) of successive images, which then decodes into a geodesic of the observation space. Orange line is the population-average trajectory, travelled across time. Blue patients has 4 visits and violet patient has 3, displaying a progression of brain alterations.

3. Conclusion

We provided a geometric interpretation the LVAE model that illustrates how the neural network can be understood as a way to parametrize the Riemannian metric of the observation manifold. The LVAE thus provides an extension of existing well-validated geometric models and provides a cheap way to compute nonlinear dynamics for high dimensional neuroimaging data.

References

Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent space oddity: on the curvature of deep generative models. *arXiv preprint arXiv:1710.11379*, 2017.

Nutan Chen, Alexej Klushyn, Richard Kurle, Xueyan Jiang, Justin Bayer, and Patrick Smagt. Metrics for deep generative models. In *International Conference on Artificial Intelligence and Statistics*, pages 1540–1550. PMLR, 2018.

Luca Falorsi, Pim De Haan, Tim R Davidson, Nicola De Cao, Maurice Weiler, Patrick Forré, and Taco S Cohen. Explorations in homeomorphic variational auto-encoding. *arXiv preprint arXiv:1807.04689*, 2018.

Samuel Gruffaz, Pierre-Emmanuel Poulet, Etienne Maheux, Bruno Jedynak, and Stanley Durrleman. Learning riemannian metric for disease progression modeling. *Advances in Neural Information Processing Systems*, 34, 2021.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

I. Koval, J-B Schiratti, A. Routier, M. Bacci, O. Colliot, S. Allassonnière, and S. Durrleman. Statistical learning of spatiotemporal patterns from longitudinal manifold-valued networks. In *International Conference on MICCAI*, pages 451–459. Springer, 2017.

Maxime Louis, Raphaël Couronné, Igor Koval, Benjamin Charlier, and Stanley Durrleman. Riemannian geometry learning for disease progression modelling. In *International Conference on Information Processing in Medical Imaging*, pages 542–553. Springer, 2019.

Benoît Sauty and Stanley Durrleman. Progression models for imaging data with longitudinal variational auto encoders. In *MICCAI 2022, International Conference on Medical Image Computing and Computer Assisted Intervention*, 2022a.

Benoît Sauty and Stanley Durrleman. Riemannian metric learning for progression modeling of longitudinal datasets. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2022b.

J-B Schiratti, S. Allassonniere, O. Colliot, and S. Durrleman. Learning spatiotemporal trajectories from manifold-valued longitudinal data. In *Neural Information Processing Systems*, number 28, 2015.

J-B Schiratti, S. Allassonnière, O. Colliot, and S. Durrleman. A bayesian mixed-effects model to learn trajectories of changes from repeated manifold-valued observations. *The Journal of Machine Learning Research*, 18(1):4840–4872, 2017.

Hang Shao, Abhishek Kumar, and P Thomas Fletcher. The riemannian geometry of deep generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 315–323, 2018.