Predicting structural brain trajectories with discrete optimal transport normalizing flows

Mireia Masias
Universitat Pompeu Fabra
mireia.masias@upf.edu

Miguel Ángel González Ballester
Universitat Pompeu Fabra,
Institució Catalana de Recerca i Estudis Avançats (ICREA)
ma.gonzalez@upf.edu

Gemma Piella
Universitat Pompeu Fabra
gemma.piella@upf.edu

Abstract

We propose to use discrete optimal transport normalizing flows (OT-NF) for the simultaneous synthesis of brain images through the years, and the explicit control of such progression. The OT-NF formulation, based on the minimization of the sliced-Wasserstein distance, allows inferring such trajectories in the absence of longitudinal data. The proposed framework could allow the imputation of brain images, conditioned on non-constant insults or stimuli.

1 Introduction

From birth to death, the brain undergoes subtle and non-linear changes. Extensive clinical research has been conducted to characterize them throughout their lifetime. These studies are based on the recruitment and analysis of brain images at different time points. Some of them have succeeded in following up on the same subjects through the years. However, the number of subjects undergoing all the acquisitions is usually very low, requiring a great deal of effort and a considerable amount of time to obtain complete enough databases. Consequently, most of these studies have opted to infer population-level trends using regression models or atlases that best fit cross-sectional data at different time points.

Thanks to them, we now know how the average brain develops and degenerates, and we can approximately predict how a specific brain will progress over the years. Nevertheless, changes might be heterogeneous in the population, both spatially and temporally, traits that the previous methods might not be able to capture. There is therefore a need to find more precise predictive methods. This could have a range of clinical implications, from the imputation of longitudinal data in incomplete datasets to the detection of small variations with respect to the expected trajectories, which could allow early diagnosis of neurodevelopmental or neurodegenerative conditions.

The use of generative models may be suitable for the described problem, as they have been proven to be very effective in the generation of high-quality synthetic data. To date, few studies have proposed similar approaches for brain prognosis. On the one hand, [1] trained a conditional variational autoencoder (cVAE) to predict the metabolic topography of subjects at different ages, whereas [7] built a generative adversarial network (GAN) that was able to predict morphological age progression while preserving subject's identity. In both cases, networks could be conditioned on additional clinical information and were trained using cross-sectional data.

Although providing very accurate results, none of these methods offer explicit control over the trajectory of latent variables. The manifold on which predicted system dynamics are embedded

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

remains unknown and conditioning on additional clinical information, by design, can only be introduced as a constant factor. Recently, [6] and [2] proposed to consider the challenge of estimating these paths as one of dynamic optimal transport. To do so, the authors proposed the use of normalizing flows (NF). In their original formulation, NF are generative models able to learn bijective functions approximating any complex distribution from a simpler known distribution, such as a Gaussian, taking advantage of the change of variable rule; but the authors expanded the method allowing the network to learn such transforms between any pair of unknown distributions, using the minimization of the Wasserstein distance, providing promising results in the continuous and discrete formulations, respectively.

In this work, we exploit this idea to explicitly model individual brain trajectories, using cross-sectional data in a reduced space, while allowing to get back individual predicted synthetic images at different ages.

2 The proposed method

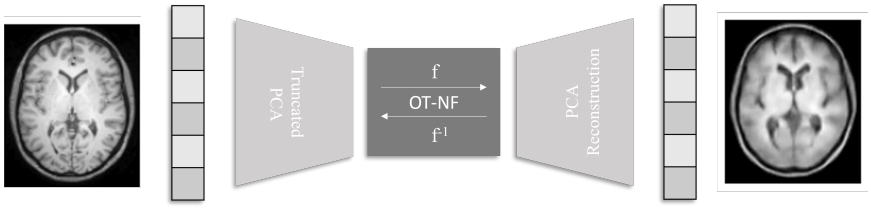

Figure 1: General algorithmic flow.

The dimensionality of input images was first reduced using principal component analysis (PCA). The first 50 components were used, keeping 70% of the total variance. The discrete optimal transport normalizing flow (OT-NF) was implemented in this low-dimensional space.

For the NF, we particularly implemented the RealNVP architecture [3], composed of 6 coupling layers (CL), each of them encoding for a specific age range between 20 and 70 years old or more, which was the range included in the database used. Thus, each of these CL constituted not only the transformation core of the network but also the estimator for each of the age ranges.

Each CL was composed of 3 steps. For each step, half of the latent variables were masked (x_i^1), whereas the other half (x_i^2) was used for the estimation of scaling (s) and translation (t) factors. These factors were then used for the affine transformation of x_i^1:

$$\begin{aligned} x_{i+1}^2 &= x_i^2 \\ x_{i+1}^1 &= s(x_i^2)x_i^1 + t(x_i^2) \end{aligned} \tag{1}$$

Where i stands for the transformation step. In this manner, the transformation is invertible and the Jacobian of the transform is triangular, with an easy-to-compute determinant. For each step in the CL, the masking was alternated. The factors s and t were computed through the training of small networks composed of 5 dense layers, with *ReLU* activation function, except for the last layer with *linear* and *tanh* activation functions, for t and s respectively.

To train the discrete OT-NF, here we adopted the strategy proposed by [2]. Monge's formulation of OT states that given a pair of probability measures μ and ν, from X and Y metric spaces, there exists a diffeomorphism $T : X \to Y$ such that it allows reaching an infimum for a given cost function c:

$$\inf_T \left\{ \int_X c(x, T(x)) d\mu(x) \middle| T^*(\mu) = \nu \right\} \tag{2}$$

Being $T^*(\cdot)$ the forward operator of $T(\cdot)$. The latter diffeomorphism is called the optimal transportation map. However, the existence and uniqueness of such a function cannot be guaranteed. Therefore, the authors propose a relaxation of Monge's OT problem, by replacing the equality $T^*(\mu) = \nu$ with the minimization of the distance between $T^*(\mu)$ and ν. In this manner, given a cost function c, now the OT problem, in the discrete form, can be relaxed to the following training loss:

$$\min_{T}\left\{ d(T^{*}(\mu),\nu) + \lambda \sum c(x, T(x)) \right\} \tag{3}$$

In this case, we have chosen $d(\cdot,\cdot)$ to be the sliced-Wasserstein distance (SWD) between the output of the last CL and the target distribution. The Euclidean distance between adjacent transformed points was chosen to be the regularizer c, and λ was set to 1. Since the shortest Euclidean distance between two points is a straight line, in a first attempt to preserve the manifold of the transformations, we added an extra loss consisting of the SWD between the output of some of the CL and the corresponding experimental age distribution.

For the dimensionality reduction, the scikit-learn implementation of the PCA (v. 1.1.1.) was employed. The discrete OT-NF was implemented using the Keras library (v.2.4.0), and the network was trained using a gradient descent algorithm (Adam [5]).

3 Preliminary results

T1-weighted MR images from the IXI-Database [4] were used, a dataset designed for the study of brain development. It contains images from healthy subjects in an age range between 20 and 80 years old. The subsample used was acquired at two different scanners: a 1.5T Philips system (332 subjects from Guy's Hospital, London, UK), and a 3T Philips system (185 subjects from Hammersmith Hospital, London, UK). All images were normalized, bias-field corrected, and transformed to the MNI space, using both rigid and affine transforms. For each subject, 40 consecutive centered axial slices were included. The selected dataset was reduced using PCA, and the first two PCs were used to plot distributions for monitoring purposes.

Once trained, the discrete OT-NF was used to transform the input distribution, in black (i.e., brain images at age twenty); to the target distribution, in red (i.e., brain images at age seventy), as seen in Figure 2. The output of each CL, also in black, represents the estimated distribution for the corresponding age range. By performing the inverse PCA on the transformed latent variables, we obtained an estimate of specific brain images in each of the corresponding ages.

4 Conclusion and limitations

We propose an approach for structural brain image prediction with explicit control of sample trajectories in the absence of longitudinal data for training, taking advantage of discrete OT-NF. We present here some results, in which OT paths were estimated in a low-dimensional space obtained by PCA. As a consequence, the quality of the retrieved predicted images was not yet comparable to those of [1] or [7], and only general aging trends could be discerned in the resulting images, the most notable of which were the decrease in brain tissue with the associated increase in ventricle volume. Attention to detail is of paramount importance for clinical applications. Therefore, future work includes improving image quality. One possibility is the use of other more powerful dimensionality reduction techniques, such as VAE and its variants; while another possibility is a multiscale approach, taking advantage of alternative computationally efficient but accurate SWD metrics. Moreover, the results need to be validated using real longitudinal data sets. Despite all the aforementioned, the proposed framework allows synthesizing brain images while simultaneously monitoring brain progression, opening the possibility of conditioning by insults/stimuli affecting only specific time points.

5 Aknowledgements

This publication is part of the project PCI2021-122044-2A, funded by the project ERA-NET NEURON Cofund2, by MCIN/AEI/10.13039/501100011033/ and by the European Union "NextGenerationEU"/PRTR. G. Piella is supported by ICREA under the ICREA Academia programme.

6 Potential societal impact

In this abstract, we propose the use of discrete OT-NF for the prediction of brain progression. Such application could have a range of clinical implications, from the imputation of data in incomplete

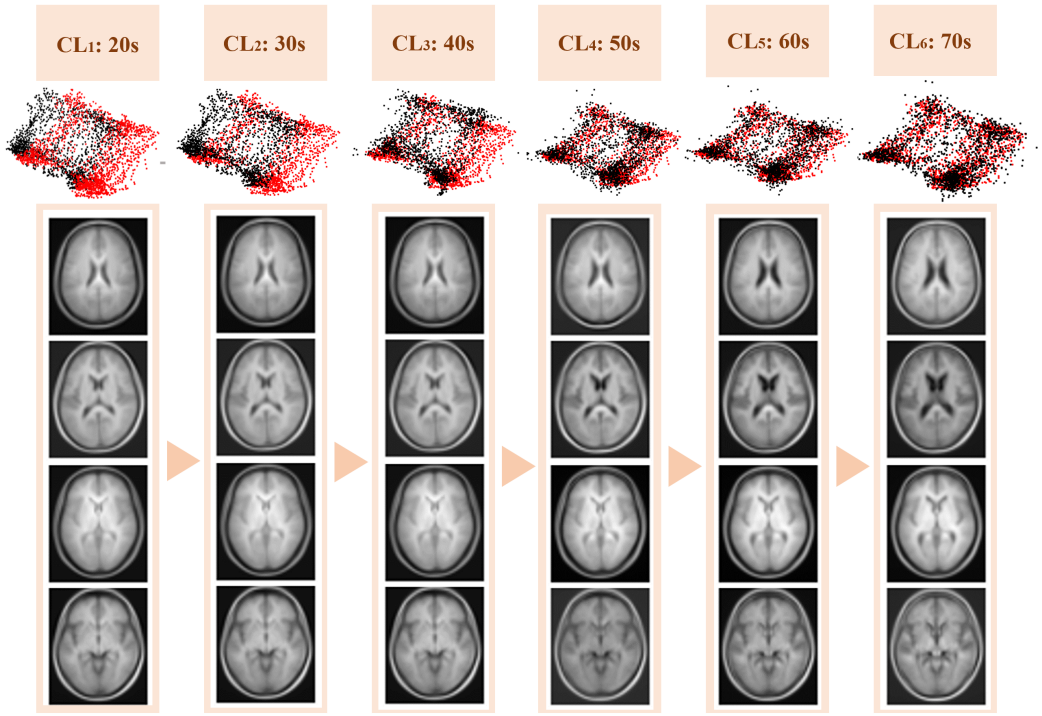

Figure 2: Detail on the discrete OT-NF architecture, with the changing distribution, represented using the first two PCs; and the corresponding predicted brain images per decade/CL.

datasets to the early detection of deviations from expected trajectories in neurodevelopmental and neurodegenerative disorders. Nevertheless, attention to detail and extensive validation of the method is required before reaching any potential clinical use. Moreover, in any of the cases, such methods should not be designed for replacing medical knowledge, but as auxiliary tools for decision-making. Therefore, interpretability is of utmost importance. In this direction, the formulation of the proposed method should have an additional advantage when compared to other approaches, as it allows the explicit modeling of the manifold of trajectories.

References

[1] Hongyoon Choi, Hyejin Kang, and Dong Soo Lee. Predicting aging of brain metabolic topography using variational autoencoder. *Frontiers in Aging Neuroscience*, 10, 7 2018.

[2] Florentin Coeurdoux, Nicolas Dobigeon, and Pierre Chainais. Learning optimal transport between two empirical distributions with normalizing flows. 7 2022.

[3] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real-nvp. 5 2016.

[4] DL Hill, S Williams, D Hawkes, and S.M. Smith. Open dataset: Ixi - information extraction from images project (epsrc gr/s21533/02)., 2006.

[5] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 12 2014.

[6] Alexander Tong, Jessie Huang, Guy Wolf, David van Dijk, and Smita Krishnaswamy. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics. 2 2020.

[7] Tian Xia, Agisilaos Chartsias, Chengjia Wang, and Sotirios A. Tsaftaris. Learning to synthesise the ageing brain without longitudinal data. *Medical Image Analysis*, 73:102169, 10 2021.