Metric regularization of latent spaces via Ricci-type flows

16/11/2023

Alexey Lazarev, PhD student at IMT, Toulouse.
PhD advisors: Reda Chhaibi and François Costantino.

ANITI

Université de Toulouse

Latent space and its metric

Autoencoder structure

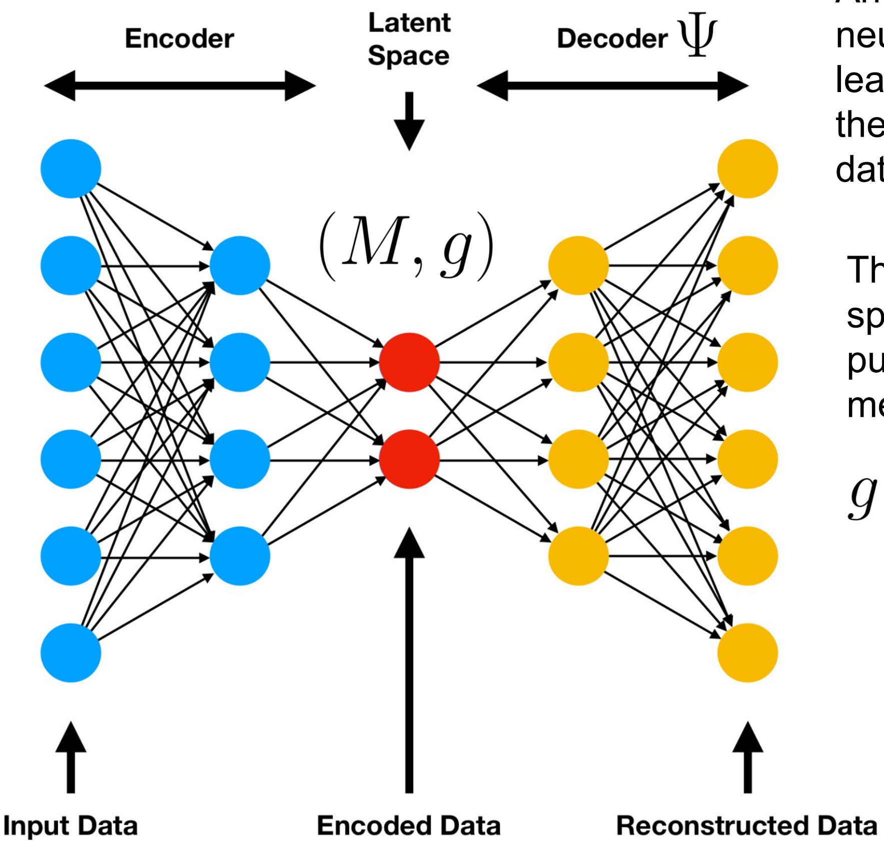

An autoencoder is a neural network that learns to compress and then reconstruct input data.

The metric in the latent space is the Riemannian pullback of the Euclidean metric:

$$
\begin{aligned}
g &= \Psi^* \| \cdot \|_2^2 \\
&= \nabla \Psi^* \nabla \Psi.
\end{aligned}
$$

Unorganized latent space

Well organized latent space

Ricci flow as a gradient flow, following Perelman

Ricci flow on manifolds

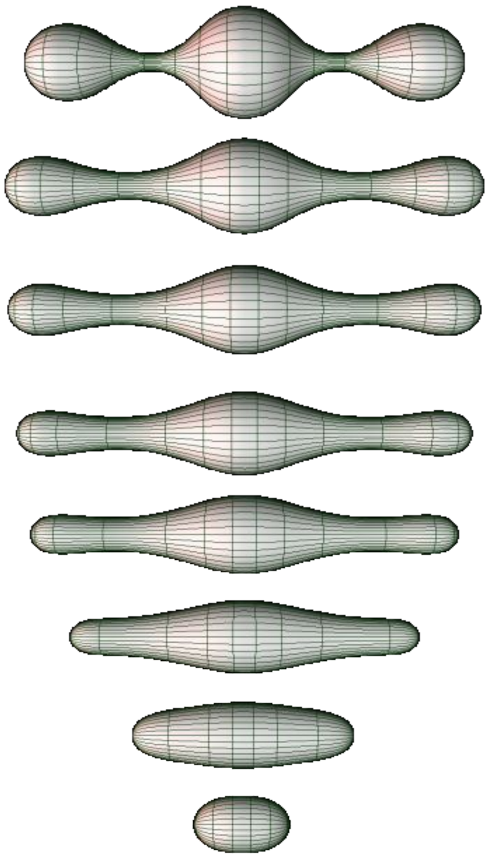

Loss function: $\mathcal{L} = \lambda_{\text{recon}} \text{MSE} + \lambda_{\text{curv}} \mathcal{L}_{\text{curv}},$

where $\mathcal{L}_{\text{curv}}$ is any of the functionals:

$$\mathcal{L}_{\text{curv}}(g) = \int_M R^2 d\mu, \quad (1)$$

$$\mathcal{F}(g, f) = \int_M (R + |\nabla f|^2) e^{-f} d\mu, \quad (2)$$

$$\mathcal{W}(g, f, \tau) = \int_M \frac{\tau(|\nabla f|^2 + R) + f - n}{(4\pi\tau)^{\frac{n}{2}} e^f} d\mu. \quad (3)$$

f is a potential whose gradient drives a diffeomorphism flow,
τ is a time factor. Both give a gauge transformation relating the gradient flow of $\mathcal{W}$ to the true Ricci flow.

In fact, $R = F_{\text{computable}}(\nabla\Psi, \nabla^2\Psi, \nabla^3\Psi)$ with Ψ decoder.

We use the Autograd package for automatic differentiation and computing higher order derivatives.

Datasets

1) Synthetic Gaussians dataset:
3 Gaussians in 2d planes randomly embedded into higher dimension D.

2) Swiss roll dataset:

3D representations of the datasets

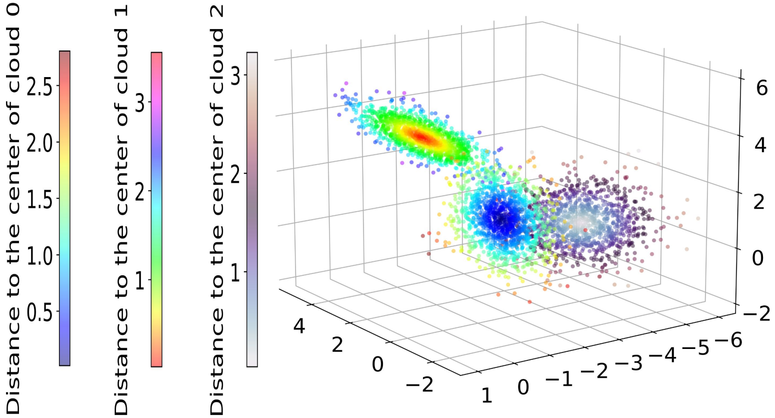

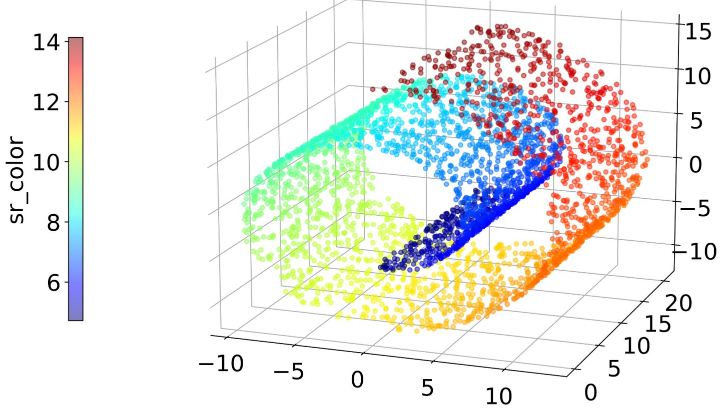

Goal: realign the Gaussian clouds on the same 2d plane, without distortion.

Goal: unroll the “doe” without tearing.

General goal: Nicer latent spaces

3D representation | t-SNE embedding | AE latent space | Ideal 2D representation

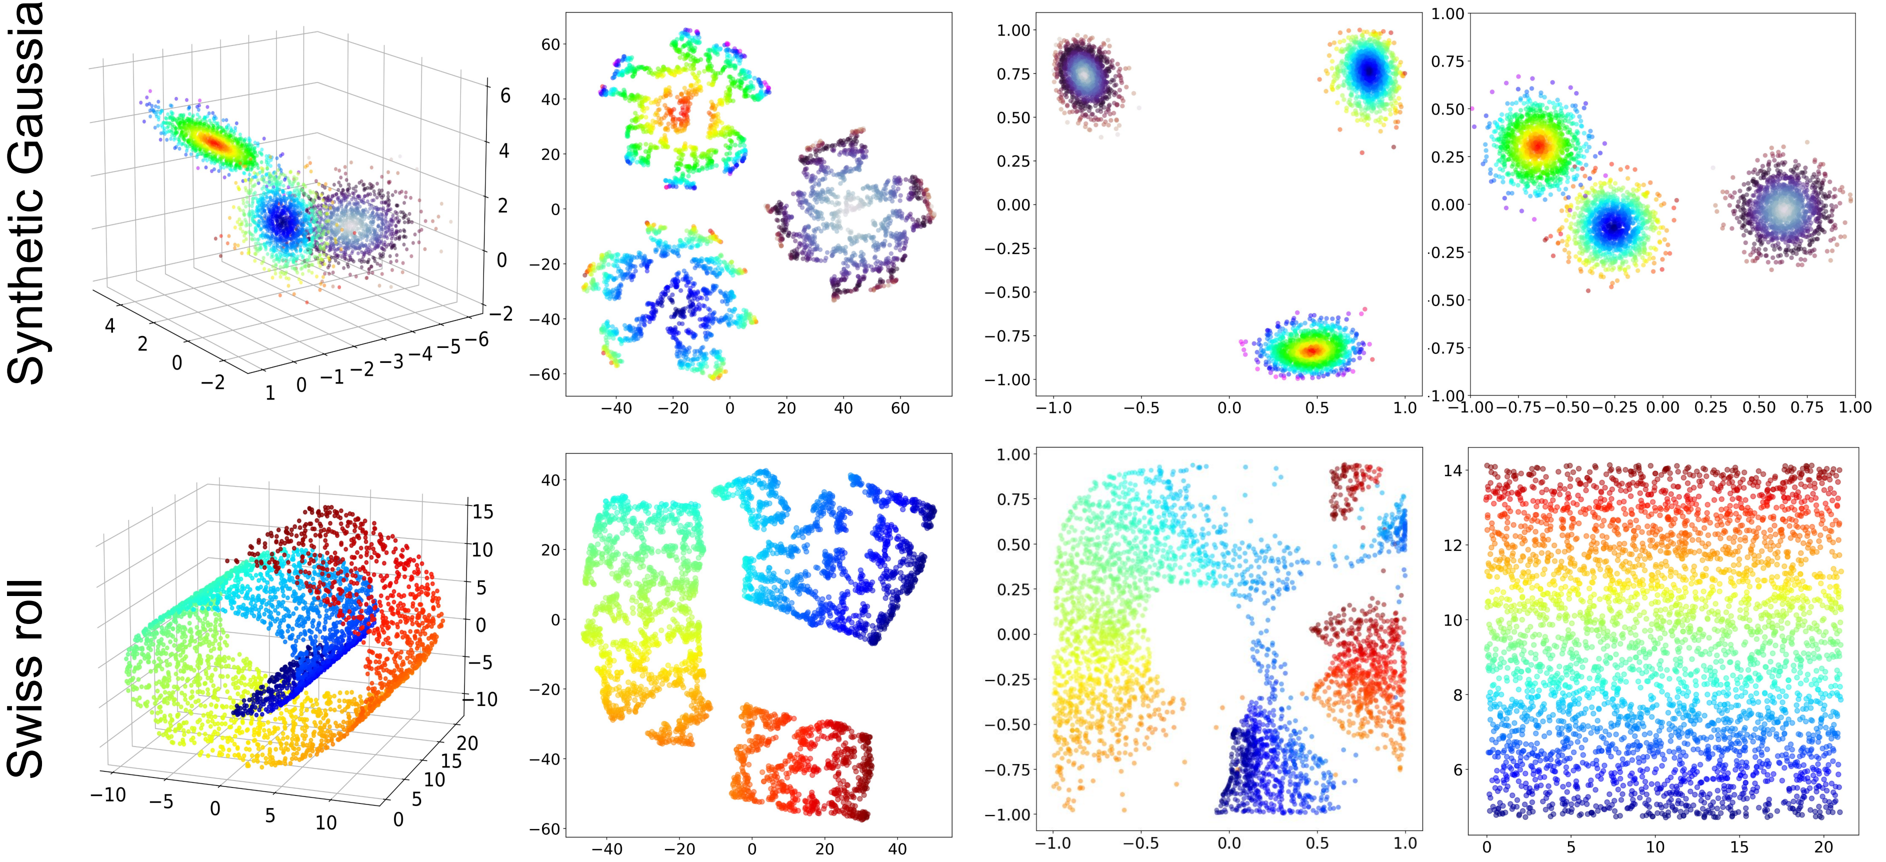

Thank you for your attention!