



HAL
open science

Data Generation in Low Sample Size Setting Using Manifold Sampling and a Geometry-Aware VAE

Clément Chadebec, Stéphanie Allasonnière

► **To cite this version:**

Clément Chadebec, Stéphanie Allasonnière. Data Generation in Low Sample Size Setting Using Manifold Sampling and a Geometry-Aware VAE. 2021. hal-03180775v1

HAL Id: hal-03180775

<https://hal.science/hal-03180775v1>

Preprint submitted on 25 Mar 2021 (v1), last revised 21 Oct 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Data Generation in Low Sample Size Setting Using Manifold Sampling and a Geometry-Aware VAE

Clément Chadebec¹ Stéphanie Allasonnière¹

Abstract

While much efforts have been focused on improving Variational Autoencoders through richer posterior and prior distributions, little interest was shown in amending the way we generate the data. In this paper, we develop two non *prior-dependent* generation procedures based on the geometry of the latent space seen as a Riemannian manifold. The first one consists in sampling along geodesic paths which is a natural way to explore the latent space while the second one consists in sampling from the inverse of the metric volume element which is easier to use in practice. Both methods are then compared to *prior-based* methods on various data sets and appear well suited for a limited data regime. Finally, the latter method is used to perform data augmentation in a small sample size setting and is validated across various standard and *real-life* data sets. In particular, this scheme allows to greatly improve classification results on the OASIS database where balanced accuracy jumps from 80.7% for a classifier trained with the raw data to 89.1% when trained only with the synthetic data generated by our method. Such results were also observed on 4 standard data sets.

1. Introduction

Despite the apparent availability of always bigger data sets, the lack of data remains a key issue for many fields of application. One of them is medicine where practitioners have to deal with potentially very high dimensional data (*e.g.* functional Magnetic Resonance Imaging for neuroimaging) along with very low sample sizes (*e.g.* rare diseases) which make statistical analysis challenging and unreliable (Button et al., 2013; Turner et al., 2018). In addition, the strong development of algorithms heavily relying on the deep learn-

ing framework (Goodfellow et al., 2016) and requiring a large amount of data has made the need for data augmentation (DA) crucial to avoid poor performance or over-fitting (Shorten & Khoshgoftaar, 2019). As an example, a classic way to perform DA on images consists in applying simple transformations such as random noise or cropping. However, it may be easily understood that such augmentation techniques are strongly data set dependant¹ and may still require *prior* knowledge of the data or the intervention of an expert assessing the relevance of the augmented samples. The recent development of generative models such as Generative Adversarial Networks (GAN) (Goodfellow et al., 2014) or Variational AutoEncoders (VAE) (Kingma & Welling, 2014; Rezende et al., 2014) paves the way for consideration of another way to augment the training data. While GANs have already seen some success (Calimeri et al., 2017; Zhu et al., 2017; Lim et al., 2018; Antoniou et al., 2018) and even for medical data (Liu et al., 2019; Sandfort et al., 2019) VAEs have been of less interest or involve marginal gains (Zhuang et al., 2019). We can nonetheless cite the works of (Nishizaki, 2017; Hsu et al., 2017; Wu et al., 2019) who used VAE for acoustic models.

One limitation of the use of such generative models relies in their need of a large amount of data to be able to generate faithfully. Moreover, the generative capacity of vanilla VAE has been widely discussed since when compared to GANs, VAEs produce most of the time fuzziier samples. In recent years much effort have been focused on improving the model by enhancing the variational approximate distribution using for example MCMC steps (Salimans et al., 2015) or *normalizing flows* (Rezende & Mohamed, 2015) aiming at improving the Evidence Lower Bound (ELBO) (Alemi et al., 2016; Burda et al., 2016; Higgins et al., 2017; Zhang et al., 2018). While these steps resulted in major improvements of the model, it was also noted that as expected the prior distribution plays a central role as well (Hoffman & Johnson, 2016). A natural improvement of the classic standard Gaussian prior consists in considering a Mixture of Gaussian (Nalisnick et al., 2016; Dilokthanakul et al., 2017) which was further improved with the Variational Mixture of Posterior (Tomczak & Welling, 2018). Noting that the

¹Université de Paris, INRIA, INSERM, Centre de Recherche des Cordeliers, UMRS 1138, Sorbonne Université. Correspondence to: Clément Chadebec <clement.chadebec@sorbonne-universite.fr>.

¹Think of digits where rotating a 6 gives a 9 for example.

vanilla VAE fails to apprehend data with a specific geometry, modellings of the latent space as a hypersphere (Davidson et al., 2018) or a Riemannian manifold (Arvanitidis et al., 2018; Chen et al., 2018; Mathieu et al., 2019; Kalatzis et al., 2020) were also proposed. These geometric considerations recently led to the proposal of a model combining Riemannian Hamiltonian dynamics (Girolami et al., 2009; Girolami & Calderhead, 2011) and metric learning allowing to learn a Riemannian structure over the latent space (Chadebec et al., 2020). Nevertheless, the authors still used a standard normal distribution to generate data which we believe allows for a poor latent space prospecting. While changing the prior or the posterior distribution has been widely explored, amending the whole generation procedure has still been poorly discussed. One example may be nonetheless be found in (Arulkumaran et al., 2016) where the authors proposed to use a Gibbs sampler like generation procedure but which produces correlated samples.

In this paper, we take a rather different approach and argue that using the prior distribution to generate new samples is often sub-optimal. Firstly, it does not take into account the whole structure of the latent space. Secondly, since the choice of the prior is data dependent and is conditioned by the tractability of the ELBO, it is rather hard to find the best suited prior. Finally, the *prior-based* generation methods failure is even more visible when small data sets are considered since priors may allow the sampling of areas of the latent space with potentially very poor information. These observations led us to propose the following contributions:

- We propose a non *prior-dependent* generation method using the learned geometry of the latent space. It consists in exploring the latent space using a Riemannian random walk which aims at sampling along geodesics.
- While this is a natural way to explore the manifold, we also introduce a second method consisting in sampling from the inverse of the metric volume element which is easier to use in practice. We briefly discuss how these two methods are related.
- Finally, we use the last method to perform data augmentation in a small sample size setting and validate its robustness on standard data sets and real data from the OASIS database (Marcus et al., 2007) where it allows to remarkably improve classification results.

2. Variational Autoencoder

2.1. Model Setting

Given a set of data $x \in \mathcal{X}$, the objective of a VAE is to maximise the likelihood of the associated parametric model $\{\mathbb{P}_\theta, \theta \in \Theta\}$. Assuming that there exist continuous latent variables $z \in \mathcal{Z}$ living in a smaller space \mathcal{Z} , the marginal

distribution writes

$$p_\theta(x) = \int_{\mathcal{Z}} p_\theta(x|z)q(z)dz, \quad (1)$$

where q is a prior distribution over the latent variables acting as a regularisation factor and $p_\theta(x|z)$ is most of the time a simple distribution and is referred to as the *decoder*. Since the integral of Eq. (1) is most of the time intractable so is the posterior distribution. To address such an issue, the idea is to use Variational Inference and introduce a variational distribution q_φ (often taken as Gaussian) aiming at approximating the true posterior distribution and is referred to as the *encoder* (Kingma & Welling, 2014). Using Importance Sampling allows to derive an unbiased estimate of $p_\theta(x)$

$$\hat{p}_\theta(x) = \frac{p_\theta(x|z)q(z)}{q_\varphi(z|x)} \quad \text{and} \quad \mathbb{E}_{z \sim q_\varphi} [\hat{p}_\theta] = p_\theta(x).$$

Therefore, a lower bound on the logarithm of the objective function of Eq. (1) can be derived using Jensen’s inequality:

$$\log p_\theta(x) \geq \underbrace{\mathbb{E}_{z \sim q_\varphi} [\log p_\theta(x, z) - \log q_\varphi(z|x)]}_{ELBO}. \quad (2)$$

Using the reparametrization trick makes the ELBO tractable and so can be optimised with respect to both θ and φ , the *encoder* and *decoder* parameters.

2.2. Enhancing the Model

Since the ELBO is only a lower bound on the true objective function, several methods to make it tighter were proposed. Trying to tweak the approximate posterior distribution is one of them. To do so, parametrized invertible mappings called *normalizing flows* (NF) were proposed to *sample* z and led to richer distributions (Rezende & Mohamed, 2015). Ideally, we would like to have access to NFs targeting the true posterior. In that respect, a model inspired by the Hamiltonian Monte Carlo sampler (HMC) (Neal & others, 2011) was proposed in (Salimans et al., 2015) and (Caterini et al., 2018). The strength of such a model relies in the choice of the NFs which are guided by the gradient of the true posterior. The authors of (Chadebec et al., 2020) went further and tried to enrich this approach with some geometrical considerations. In their paper, they saw the latent space as a Riemannian manifold endowed with a Riemannian metric and combined NFs and metric learning using the Riemannian Hamiltonian Monte Carlo sampler (Girolami et al., 2009; Girolami & Calderhead, 2011) to propose the Riemannian Hamiltonian VAE (RHVAE). While several papers have been trying to propose Riemannian metrics (Arvanitidis et al., 2018; Chen et al., 2018), the authors proposed to parametrized it and learn it directly from the data. Their parametrization writes

$$\mathbf{G}^{-1}(z) = \sum_{i=1}^N L_{\psi_i} L_{\psi_i}^\top \exp\left(-\frac{\|z - c_i\|_2^2}{T^2}\right) + \lambda I_d, \quad (3)$$

where N is the number of observations, L_{ψ_i} are lower triangular matrices parametrized with neural networks, c_i are the *centroids* and T (resp. λ) are non-negative hyper-parameters. While their method seems quite suited for the small sample size setting, they still used the classic generation procedure consisting in drawing a random variable from the prior and feeding it to the decoder to generate new samples. We believe that such an approach remains sub-optimal. Building upon their work, we propose two new ways of generating samples relying on the learned geometry of the latent space. First, we recall some elements of Riemannian geometry.

3. Some Elements on Riemannian Geometry

In the framework of differential geometry, one may define a Riemannian manifold \mathcal{M} as a smooth manifold endowed with a Riemannian metric g . The metric is a smooth inner product $g : p \rightarrow \langle \cdot | \cdot \rangle_p$ on the tangent space $T_p\mathcal{M}$ defined at each point p of the manifold \mathcal{M} . We call a chart x a homeomorphism mapping an open set U of the manifold to an open set V of an Euclidean space. Given $p \in U$, a chart $x : (x^1, \dots, x^d)$ induces a basis $(\frac{\partial}{\partial x^1}, \dots, \frac{\partial}{\partial x^d})_p$ on the tangent space $T_p\mathcal{M}$. Hence, a local representation of the Riemannian metric in the chart x can be written as a positive definite matrix $\mathbf{G}(p) = (g_{i,j})_{p, 0 \leq i, j \leq d} = (\langle \frac{\partial}{\partial x^i} | \frac{\partial}{\partial x^j} \rangle_p)_{0 \leq i, j \leq d}$ for $p \in \mathcal{M}$. That is for $v, w \in T_p\mathcal{M}$, we have $\langle v | w \rangle_p = v^\top \mathbf{G}(p) w$. The length of a curve γ between two points of the manifold $z_1, z_2 \in \mathcal{M}$ and parametrized by $t \in [0, 1]$ such that $\gamma(0) = z_1$ and $\gamma(1) = z_2$ is given by

$$\mathcal{L}(\gamma) = \int_0^1 \|\dot{\gamma}(t)\|_{\gamma(t)} dt = \int_0^1 \sqrt{\langle \dot{\gamma}(t) | \dot{\gamma}(t) \rangle_{\gamma(t)}} dt.$$

Curves minimising such a length are called *geodesics* and a distance on the manifold can be introduced:

$$d(z_1, z_2) = \min_{\gamma} \mathcal{L}(\gamma) \quad \text{s.t.} \quad \gamma(0) = z_1, \gamma(1) = z_2. \quad (4)$$

The manifold \mathcal{M} is said to be *geodesically complete* if all geodesic curves can be extended to \mathbb{R} . In other words at any point $p \in \mathcal{M}$, one may draw a *straight* line (with respect to the formerly defined distance) indefinitely and in any direction. For any $p \in \mathcal{M}$, the exponential map at p maps a vector v of the tangent space $T_p\mathcal{M}$ to a point of the manifold $\tilde{p} \in \mathcal{M}$ such that the geodesic starting at p with initial velocity v reaches \tilde{p} at time 1.

$$\text{Exp}_p : \begin{cases} T_p\mathcal{M} & \rightarrow \mathcal{M} \\ v & \rightarrow \text{Exp}_p(v) = \gamma_{(p,v)}(1) \end{cases},$$

where $\gamma_{(p,v)}$ is such that $\gamma(0) = p$ and $\dot{\gamma}(0) = v$. In particular, if the manifold is *geodesically complete*, then Exp_p is defined on the entire tangent space $T_p\mathcal{M}$ for any $p \in \mathcal{M}$.

Conversely, the inverse of Exp_p is Log_p which maps a point of the manifold \mathcal{M} to the tangent space $T_p\mathcal{M}$. A statistical framework can be developed on such manifolds allowing to define a measure $d\mathcal{M}(p) = \sqrt{\det g(p)} dp$ and so probability densities on Riemannian manifolds (Pennec, 2006).

4. The Proposed Methods

In this paper, we propose two new sampling procedures exploiting the Riemannian structure of the latent space and independent from the choice of the prior distribution.

4.1. Setting

The view we adopt is to consider the VAE as a tool to perform dimensionality reduction by extracting the latent structure of the data within a smaller space. Having learned such a structure, we propose to exploit it to enhance the data generation process. This differs from the standard fully probabilistic view which uses the prior distribution to generate. We believe that this is far from being optimal since the prior appears quite strongly data dependent and must remain simple so that the ELBO is tractable. We will adopt the same setting as the one exposed in (Chadebec et al., 2020). The latent space is the Riemannian manifold (\mathbb{R}^d, g) where g is the Riemannian metric whose local representation is given by Eq. (3). Since we work in an ambient-like manifold \mathbb{R}^d , there exists a global chart given by $\phi = id$ and we assume that we work in this coordinate system.

Proposition 1. *The Riemannian manifold (\mathbb{R}^d, g) is geodesically complete.*

Proof. See Appendix. A. □

Let's assume that we have learned a latent space \mathcal{Z} and the metric \mathbf{G} with a RHVAE. Thanks to Prop. 1, we are now able to refer to probability distributions on the manifold \mathcal{Z} .

4.2. The Wrapped Normal Distribution

The notion of normal distribution may be extended to Riemannian manifolds in several ways. First, it can be characterised as the maximum entropy distribution and involves the geodesic distance (Pennec, 2006). Another way to define a normal distribution on a manifold is to consider the so-called *wrapped normal* distribution. The main idea is to define a classic Euclidean normal distribution $\mathcal{N}(0, \Sigma)$ on the tangent space $T_p\mathcal{M}$ for any $p \in \mathcal{M}$ and pushing it forward to the manifold using the exponential map. This defines a probability distribution on the manifold $\mathcal{N}^W(p, \Sigma)$ called the *wrapped normal* distribution. Sampling from this distribution is straight forward and consists in drawing a velocity in the tangent space from $\mathcal{N}(0, \Sigma)$ and mapping it onto the manifold using the exponential map (Mallasto




Figure 1. Geodesic shooting using the exponential map for different starting points (red dots), initial velocities (orange arrows) and learned metrics whose log-volume element ($\sqrt{\det \mathbf{G}(z)}$) is presented in gray scale. The resulting paths are displayed in orange.




Figure 2. 500 samples from of a *standard* wrapped Gaussian distribution $\mathcal{N}^W(z, I_d)$ for different means (purple dots). The mean of the posterior associated to the training points is highlighted with the crosses. The model is trained on 180 binary circles and rings (see Appendix. D for parameters settings).

& Feragen, 2018). In this paper, we decide to work with such a distribution since the Riemannian distance of Eq. (4) involves an optimisation problem which is pretty hard to solve in general. Moreover, the *wrapped normal* allows for a latent space prospecting along geodesic paths. To be able to use it we nonetheless need to compute Exp_p which can be performed with a numerical scheme (see. Appendix. C). In Fig. 1 are displayed some geodesic paths with different learned metrics, initial velocities (orange arrows) and starting points (red dots). Fig. 2 shows 500 samples from the *standard wrapped* distribution $\mathcal{N}^W(z, I_d)$ in two different locations of the latent space. The RHVAE model is trained on a synthetic data set composed by 180 binary images of circles and rings with different diameters and thicknesses. The log of the metric volume element $\sqrt{\det \mathbf{G}}$ is presented in gray scale while the means of the posteriors associated to training data are represented by the crosses. As expected this distribution takes into account the local geometry of the manifold thanks to the geodesic shooting steps. This is a very interesting property since it encourages the sampling to remain close to the data as geodesics tend to travel through locations with the lowest volume elements.

4.3. Method 1: Riemannian Random Walk

A natural way to explore the latent space of a VAE is to use a random walk like algorithm which moves from one location to another with a certain probability. The idea here is to create a *geometry-aware* Markov Chain $(z^t)_{t \in \mathbb{N}}$ where z^{t+1} is sampled using the *wrapped* normal $z^{t+1} \sim \mathcal{N}^W(z^t, \Sigma)$. However, a drawback of such a method is that every sample of the chain is accepted regardless of its relevance. This is particularly annoying in a low sample size setting where most of the latent space has not been prospected and thus contains poor information (see Sec. 5.1).

4.3.1. QUANTIFYING THE SAMPLES QUALITY

Ideally, we would like to have access to a metric providing information on the quality of z^t . This would allow us to ignore elements of the chain potentially producing *irrelevant* samples when decoded. This measure may be provided by the Riemannian metric. By design, the learned metric is such that it has a high volume element far from the data (Chadebec et al., 2020). This implies that the volume element encodes in a way the amount of information contained at a specific location of the latent space. The higher the volume element, the less information we have. The same idea was used in (Lebanon, 2006) where the author proposed to see the inverse metric volume element as a maximum likelihood objective to perform metric learning. This aimed at finding the best metric ensuring that geodesic paths remain close to the data. In our case the likelihood definition writes

$$\mathcal{L}(z) = \frac{\rho_S(z) \sqrt{\det \mathbf{G}^{-1}(z)}}{\int_{\mathbb{R}^d} \rho_S(z) \sqrt{\det \mathbf{G}^{-1}(z)} dz}, \quad (5)$$

where $\rho_S(z) = 1$ if $z \in S$, 0 otherwise, and S is taken as a compact set so that the integral is well defined. We propose to use this measure to assess the samples quality as an *acceptance-rejection* rate α in the chain

$$\alpha(\tilde{z}, z) = \min \left(1, \frac{\sqrt{\det \mathbf{G}^{-1}(\tilde{z})}}{\sqrt{\det \mathbf{G}^{-1}(z)}} \right),$$

where z is the current state of the chain and \tilde{z} is the proposal obtained by sampling from the *wrapped* Gaussian $\mathcal{N}^W(z, \Sigma)$. The idea is to compare the relevance of the proposed sample to the current one. The ratio is such that any new sample improving the likelihood metric \mathcal{L} is automatically accepted while a sample degrading the measure is more likely to be rejected. Finally, a pseudo-code of the proposed method is provided in Alg. 1.

4.3.2. DISCUSSION

It may be easily understood that the choice of the covariance matrix Σ in Alg. 1 has quite an influence on the resulting

Algorithm 1 Riemannian random walk

Input: z_0, Σ
for $t = 1 \rightarrow T$ **do**
 Draw $v_t \sim \mathcal{N}(0, \Sigma)$
 $\tilde{z}_t \leftarrow \text{Exp}_{z_{t-1}}(v_t)$
 $\alpha = \min \left(1, \frac{\sqrt{\det \mathbf{G}^{-1}(\tilde{z}_t)}}{\sqrt{\det \mathbf{G}^{-1}(z_{t-1})}} \right)$
 Accept the proposal \tilde{z}_t with probability α
 if Accepted **then**
 $z_t \leftarrow \tilde{z}_t$
 else
 $z_t \leftarrow z_{t-1}$
 end if
end for
Return $[z_1, \dots, z_T]$

sampling. On the one hand, a Σ with strong eigenvalues will imply drawing velocities of potentially high magnitude allowing for a better prospecting. Nonetheless, the proposals are more likely to be rejected since geodesics may go farther and end in areas with lower likelihood. On the other hand, small eigenvalues involve a high acceptance rate but it will take a longer time to prospect the manifold. An adaptive method where Σ depends on \mathbf{G} may be considered and will be part of future work. Moreover, even though the computation of the exponential is pretty straightforward, it involves the use of a numerical scheme adding some computational cost to the method.

4.4. Method 2: Metric Sampling

For the second method referred to as *metric sampling*, we take a rather different approach since we propose to directly sample from the inverse of the metric volume element which is seen as a probability distribution on \mathbb{R}^d as in Eq. (5). As discussed above the inverse volume element may be seen as a likelihood measure assessing the quantity of information within a specific area of the latent space. This property implies that locations having the most information will also be the most prospected throughout the sampling. Since \mathbf{G}^{-1} has a closed form (see Eq. (3)), sampling from this distribution is pretty straight forward with classic MCMC sampling methods. For instance, this may be achieved with the HMC sampler (Neal, 2005) since the gradient of the target log-density is easy to compute.

While this method seems quite different from Method 1, they remain related through the choice of the covariance matrix Σ in Alg. 1. In particular, we have the following

Remark 1. *If Σ has small enough eigenvalues then Alg. 1 samples from the distribution of Eq. (5) defined on \mathbb{R}^d .*

Informally, it tells that if the covariance has small enough

eigenvalues, then the velocity drawn in Alg. 1 will have a small magnitude. Hence, geodesic curves may be approximated by straight lines and so the ratio becomes a Hasting-Metropolis ratio. This is discussed in Appendix. B. A code for both methods is given in the supplementary materials.

5. Experiments

In this section, we compare the samples quality using the proposed methods on various standard and *real-life* data sets. Then, we show that such a change in the generation process is well suited to perform data augmentation in particular in the low sample size setting. To stick to this setting we only consider data sets with tens or hundreds of samples.

5.1. Sampling Comparison

We first compare the sampling quality between *prior-based* methods and ours.

5.1.1. STANDARD DATA SETS

The methods are first validated on a hand-made synthetic data set composed by 180 binary images of circles and rings of different diameters and thicknesses (see training samples in Appendix. E). We then train a VAE with a normal prior, a VAE with a VAMP prior and a RHVAE until the ELBO does not improve for 50 epochs. Any relevant parameters setting is stated in Appendix. D.

Fig. 3 highlights the obtained samplings with each model using either the *prior-based* generation procedure or the ones proposed in this paper. The first row presents the learned latent space along with the means of the posteriors associated to the training data (crosses) and 100 latent space samples for each generation method (blue dots). The second row displays the corresponding decoded images. The first outcome of such a study is that sampling from the prior distribution $\mathcal{N}(0, I_d)$ leads to a poor latent space prospecting. Therefore, even with balanced training classes, we end up with a model over-representing certain elements of a given class (rings). This is even more striking with the RHVAE since it tends to stretch the resulting latent space. This effect seems nonetheless mitigated by the use of a multimodal prior such as the VAMP. However, another limitation of *prior-based* methods is that they may sample in locations of the latent space potentially containing very few information (*i.e.* where no data is available). Since the decoder appears to interpolate quite linearly, the classic scheme will generate images which mainly correspond to a superposition of samples (see an example with the red dots in Fig. 3 and the corresponding samples framed in red). Moreover, there is no way to assess a sample quality before decoding it and assessing visually its relevance. These limitations may lead to a (very) poor representation of the actual data set diversity

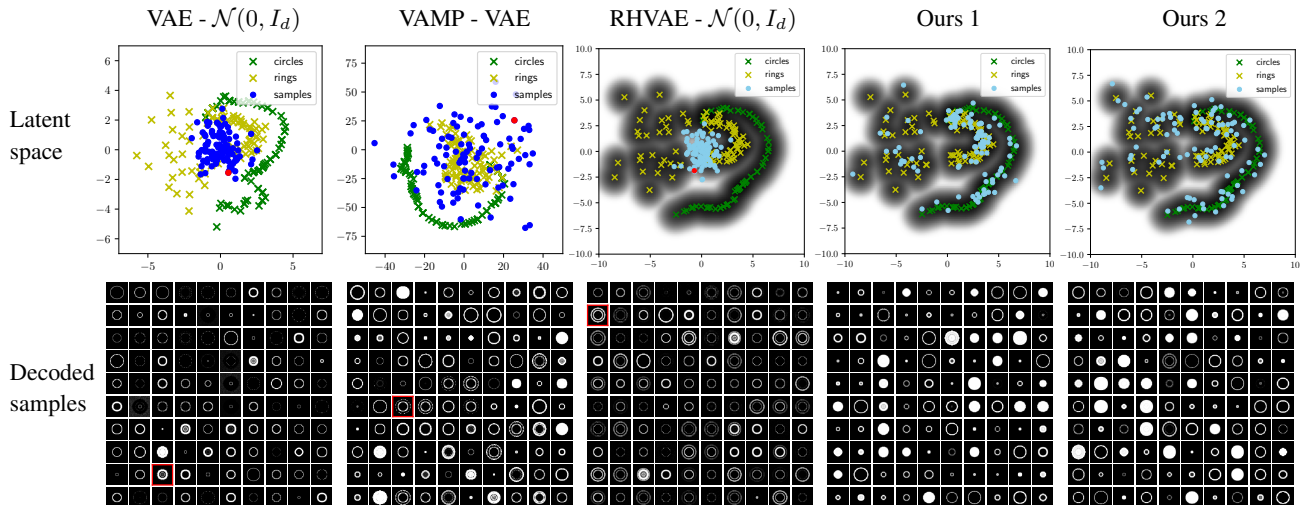


Figure 3. Comparison between prior-based generation methods and the proposed Riemannian random walk (ours 1) and metric sampling (ours 2). Top: the learned latent space with the means of the encoded training data (crosses) and 100 samples for each method (blue dots). Bottom: the resulting decoded images. The models are trained on 180 binary circles and rings with the same neural network architectures. A comparison on reduced versions of MNIST and FashionMNIST is also available in Appendix. E.

while presenting quite a few *irrelevant* samples.

Impressively, sampling along geodesic paths or using metric sampling leads to far more diverse and sharper samples. The new sampling schemes avoid regions that have been poorly prospected so that almost every decoded sample is visually satisfying and accounts for the data set diversity. Finally, neither method over-fits the training samples since 1) the samples are not always located on the means of the encoded training data, 2) it is shown in Sec. 5.2 that classifiers can achieve significantly better classification results on synthetic data than on *raw* training samples proof of generalisation power. Similar effects are observed on reduced versions of MNIST (LeCun, 1998) and FashionMNIST (Xiao et al., 2017) where samples are most of the time degraded with the classic methods while the proposed ones allow a sampling of more diverse and sharper samples (see Appendix.E).

5.1.2. OASIS DATABASE

The new generation schemes are then assessed on the publicly available OASIS database composed by 416 patients aged 18 to 96, 100 of whom have been diagnosed with very mild to moderate Alzheimer disease (AD). The MR images we use for this experiment are 208x176 gray scale images with a mask notifying brain tissues and are referred to as the *masked T88 images* in (Marcus et al., 2007). We refer the reader to their paper for further image preprocessing details. A VAE and a RHVAE are then trained to generate either cognitively normal (CN) or AD patients with the same early stopping criteria as before. Fig. 4 shows samples extracted from the training set (top), MRI generated by a vanilla VAE

(2nd row), images from the Riemannian random walk we propose (3rd row) and with the proposed metric sampling method (bottom). For each method, the upper row shows images of patients diagnosed CN while the bottom row presents an AD diagnosis. Again the proposed samplings seem able to generate a wider range of sharp samples while the VAE appears to produce non-realistic degraded images which are very similar (see red frames). For example, the proposed schemes allow us to generate realistic *old*² patients with no AD (blue frames) or younger patients with AD (orange frames) even though they are under-represented in the training set. Generating 100 images of OASIS database takes 1 min. with the first proposed method and 40 sec.³ with metric sampling with Intel Core i7 CPU (6x1.1GHz) and 16 GB RAM.

5.2. Data Augmentation

Since the proposed methods seem to be able to produce samples keeping the intrinsic topology of the data even with small size data sets, we propose to use Method 2 (*i.e.* metric sampling) which is easier to use in practice and produces apparent equivalent results to Method 1 to perform data augmentation. The idea is to 1) create a *small size* data set; 2) train a VAE to generate labelled samples for each class of the train set; 3) train a classifier on the train set, the augmented train set or the synthetic data only.

²An older person is characterised by larger ventricles.

³Depends on the chains' length (here 200 steps per image).

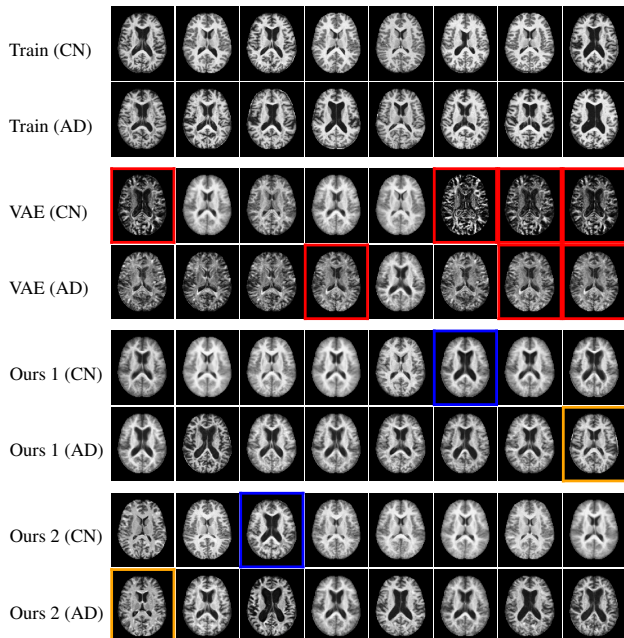


Figure 4. Generation of CN or AD patients from the OASIS database. Training samples (top), generation with a VAE and normal prior (2nd row), with the Riemannian random walk (3rd row) and with metric sampling (bottom). Generating using the prior leads to either unrealistic images or similar samples (red frames) while the proposed methods generate sharper and more diverse samples. For instance, they are able to generate CN *older* patients (blue frames) or younger AD (orange frames) even though they are under-represented within the training set.

5.2.1. STANDARD DATA SETS

First, we test the method on four reduced data sets extracted from *well-known* databases. The first data set is created by selecting 500 samples from 10 classes of MNIST ensuring balanced classes. The second one consists in selecting again 500 samples from the MNIST database but applying a random split such that some classes are *over-represented*. Then, we select 500 samples from 10 classes of the EMNIST database (Cohen et al., 2017). These classes are selected such that they are composed by both lowercase and uppercase characters so that we end up with a small database with strong variability within classes. The balance we use matches the one in the initial data set. Finally, we create another balanced one with 300 samples from 3 classes of FashionMNIST (*T-shirt*, *dress* and *shirt*). These data sets are then split such that 80% is allocated for training (referred to as the *raw data*) and 20% for validation. For a fair comparison, we use the original test set (e.g. ~ 1000 samples per class for MNIST) to test the classifiers. This ensures statistically meaningful results while assessing the generalisation power on unseen data.

5.2.2. THE AUGMENTATION METHODS

On the one hand, the *raw data* are augmented by a factor 5, 10 and 15 using classic data augmentation methods (random noise, cropping etc.). On the other hand, VAE and RHVAE models are trained individually on each class of the *raw data* until the ELBO does not improve for 20 epochs. The generative models are then used to produce 200, 500, 1000 or 2000 new synthetic samples per class using either the classic generation scheme (*i.e.* the prior) or metric sampling.

5.2.3. RESULTS

We then train a DenseNet model⁴ (Huang et al., 2017) with 5 independent runs on 1) the *raw data*; 2) the augmented data using basic transformations; 3) the augmented data using the VAE models and each generation method; 4) only the synthetic data generated by the VAEs. The DenseNet is trained until the loss does not improve on the validation set for 50 epochs. We report the mean accuracy and standard deviation on the test set for each method in Table 1.

The first outcome of such a study is that as expected generating new samples using the proposed method improves their relevance. This may be seen first by comparing the results in the second section of the table corresponding to the augmented data sets (*i.e.* *raw data* + synthetic) where metric sampling indeed allows for a quite significant gain in the model accuracy. It is even more striking when looking at the last section of Table 1 in which only synthetic samples are used to train the classifier. For instance, when we only generate 200 synthetic samples per class for MNIST with a VAE or RHVAE and train the classifier on them, the classic method fails to produce meaningful samples since a loss of 20 to 30 pts in accuracy combined with high uncertainty is observed when compared to the *raw data*. Interestingly, sampling from the inverse of the metric volume element seems to avoid such an effect. Even more impressive is the fact that we are able to produce synthetic data sets on which the classifier outperforms greatly the results observed on the *raw data* (3 to 6 pts gain in accuracy) while keeping a relatively low standard deviation (see gray cells).

Secondly, this example also shows why geometric data augmentation methods are still questionable and data set dependent. For instance, augmenting the *raw data* by a factor 10 (including flips and rotations) does not seem to have a significant effect on the MNIST data sets while it still improves results on EMNIST and FashionMNIST. Finally, the proposed method appears quite robust to data set changes and produces the best results allowing for a significant gain when compared to the *raw data*.

⁴We use the code in (Amos, 2020) (See Appendix. D).

Table 1. Data augmentation with a DenseNet model as Benchmark using either geometric transformations (Aug.), a (RH)VAE and the classic generation method or metric sampling (ours). Mean accuracy and standard deviation across 5 independent runs are reported. In gray are the cells where the accuracy is higher on synthetic data than on the *raw data*. The test set is the one proposed in the original data set (e.g. ~ 1000 samples per class for MNIST).

DATA SETS	MNIST	MNIST (UNBAL.)	EMNIST (UNBAL.)	FASHION
RAW DATA	89.89 (0.6)	81.55 (0.7)	82.61 (1.4)	76.00 (1.5)
RAW + SYNTHETIC				
AUG. (X5)	92.80 (0.4)	86.46 (0.9)	85.56 (1.3)	77.47 (2.0)
AUG. (X10)	88.25 (2.2)	81.98 (2.4)	85.75 (0.3)	79.20 (0.7)
AUG. (X15)	92.84 (0.7)	85.85 (3.4)	86.59 (0.8)	80.03 (0.5)
VAE-200*	88.53 (0.9)	84.05 (2.0)	81.74 (3.0)	78.59 (0.4)
VAE-500*	90.37 (1.4)	87.33 (1.2)	83.40 (1.6)	78.68 (0.3)
VAE-1K*	91.17 (1.0)	86.01 (2.5)	84.35 (1.6)	77.61 (2.1)
VAE-2K*	92.19 (1.6)	87.97 (2.2)	85.98 (0.2)	79.29 (1.1)
RHVAE-200*	89.90 (0.5)	82.26 (0.9)	82.96 (1.3)	77.58 (1.3)
RHVAE-500*	90.87 (1.1)	83.96 (3.2)	84.38 (1.2)	78.03 (1.3)
RHVAE-1K*	91.71 (0.8)	84.72 (1.8)	84.70 (2.4)	79.31 (1.6)
RHVAE-2K*	92.74 (1.4)	86.78 (1.0)	84.92 (2.1)	79.03 (1.4)
OURS-200*	90.98 (1.1)	84.07 (2.0)	85.11 (1.1)	76.96 (0.8)
OURS-500*	92.33 (1.1)	87.69 (0.9)	85.06 (1.1)	78.51 (0.9)
OURS-1K*	93.25 (0.8)	89.71 (0.8)	87.02 (1.0)	80.16 (0.8)
OURS-2K*	94.29 (0.8)	89.11 (1.9)	87.59 (0.8)	78.13 (1.8)
SYNTHETIC ONLY				
VAE-200*	69.89 (1.5)	64.58 (1.8)	65.67 (2.6)	73.90 (3.0)
VAE-500*	72.34 (4.2)	69.42 (4.1)	67.32 (2.4)	71.42 (8.5)
VAE-1K*	83.40 (2.4)	74.69 (3.2)	75.31 (1.4)	71.42 (6.1)
VAE-2K*	86.55 (2.2)	79.61 (3.8)	78.85 (3.0)	76.73 (1.6)
RHVAE-200*	76.03 (1.8)	61.53 (2.9)	59.84 (2.6)	72.78 (3.6)
RHVAE-500*	79.99 (2.2)	66.77 (3.3)	66.94 (4.0)	74.33 (2.6)
RHVAE-1K*	82.01 (2.9)	69.28 (1.8)	73.65 (4.1)	75.96 (4.1)
RHVAE-2K*	85.21 (3.9)	77.32 (3.2)	68.56 (2.3)	74.31 (3.1)
OURS-200*	87.19 (1.1)	79.50 (1.6)	77.05 (1.6)	76.96 (0.8)
OURS-500*	89.06 (1.3)	80.36 (2.1)	80.16 (2.0)	78.47 (0.8)
OURS-1K*	90.15 (1.4)	86.24 (1.8)	82.61 (1.3)	79.31 (0.6)
OURS-2K*	92.60 (1.1)	87.55 (1.3)	86.03 (1.0)	78.29 (0.9)

* NUMBER OF GENERATED SAMPLES

5.2.4. OASIS DATA SET

We now validate the proposed data augmentation method on a *real-life* database. We choose to use again the OASIS data set which according to us represents a nice example of day-to-day challenges practitioners have to face. We consider the binary classification problem consisting in trying to detect MRI of patients having been diagnosed AD. Then, we split the 416 images into a training set (70%) referred to as the *raw data*, a validation set (10%) and a test set (20%). The VAE and RHVAE models are trained on the *raw data* to generate either AD or CN patients until the ELBO does not improve for 50 epochs. A MLP classifier is then trained with 5 independent runs on 1) the *raw data*; 2) the augmented data using geometric transformations; 3) the augmented data using the VAEs and the classic generation or metric sampling; 5) only the synthetic data. We adopt an early stopping strategy and stop training if the validation loss does not improve for 20 epochs. The global and balanced accuracy

Table 2. Data augmentation on the OASIS database with a MLP model as Benchmark using either geometric transformations, a VAE and the classic generation method or metric sampling (ours). Global and balanced accuracy across 5 runs are reported.

DATA SETS	GLOBAL ACCURACY	BALANCED ACCURACY
RAW DATA	84.29 (2.2)	80.71 (4.0)
RAW + SYNTHETIC		
AUGMENTED (X5)	84.76 (2.3)	83.84 (3.8)
AUGMENTED (X10)	83.81 (1.8)	76.77 (3.8)
AUGMENTED (X15)	82.38 (2.0)	76.26 (3.0)
VAE-200*	81.90 (2.3)	79.59 (1.6)
VAE-500*	87.14 (1.4)	80.50 (0.5)
VAE-1000*	83.33 (0.0)	81.31 (0.0)
VAE-2000*	82.38 (0.5)	80.71 (0.3)
OURS-200*	84.76 (6.0)	85.86 (3.0)
OURS-500*	86.43 (2.7)	86.11 (1.0)
OURS-1000*	86.19 (2.3)	85.96 (2.6)
OURS-2000*	88.33 (0.5)	85.30 (1.9)
SYNTHETIC ONLY		
VAE-200*	74.52 (0.9)	80.55 (1.6)
VAE-500*	83.10 (4.8)	78.33 (1.9)
VAE-1000*	75.00 (0.0)	80.05 (0.0)
VAE-2000*	76.19 (0.0)	80.41 (0.8)
OURS-200*	89.05 (0.5)	86.97 (0.3)
OURS-500*	89.05 (0.5)	88.58 (1.1)
OURS-1000*	89.29 (0.0)	89.14 (0.0)
OURS-2000*	89.05 (0.5)	88.58 (1.1)

* NUMBER OF GENERATED SAMPLES

are available in Table. 2 (the remaining main metrics are available in Appendix. F). In this experiment, using the new generation scheme again improves overall the metrics when compared to the *raw data* and other augmentation methods. Interestingly enough, the classifier even outperforms any other method when trained on synthetic data only where mean balanced accuracy reaches 89.14 vs. 80.71 on the *raw data*. This may be due to the fact that generated samples will insist on the most dominant features of the data while removing some potentially *bias inducing* information.

6. Conclusion

In this paper, we explored two new ways of generating new samples from a VAE exclusively relying on the learned geometry of the latent space. In particular, these methods do not depend on the choice of the prior distribution which is usually used to generate. The first one consisted in defining a Riemannian random walk aiming at sampling along geodesic paths while the second one involved sampling from the inverse of the metric volume element. It was shown that those methods enhance greatly the data generation process in the low sample size setting thanks to their ability to *quantify* the amount of information contained at a specific location of the latent space. The second method, which is easy to use in practice since the inverse metric has a closed form expression, was used to perform data augmentation with well-known classifiers on several *standard*

and *real-life small* size data sets. For each experiment, the proposed method allowed for a great gain in classification results while demonstrating a strong generalisation power since the classifiers achieved better results on synthetic data than on the *raw data*. Finally, we quickly discussed how the two proposed methods are related. Future work would consist in trying to extend these generation methods to VAE models with different metrics and validate the augmentation method on even more challenging data such as volumes.

References

- Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.
- Amos, B. bamos/densenet.pytorch, December 2020. URL <https://github.com/bamos/densenet.pytorch>. original-date: 2017-02-09T15:33:23Z.
- Antoniou, A., Storkey, A., and Edwards, H. Data augmentation generative adversarial networks. *arXiv:1711.04340 [cs, stat]*, March 2018. URL <http://arxiv.org/abs/1711.04340>. arXiv: 1711.04340.
- Arulkumaran, K., Creswell, A., and Bharath, A. A. Improving sampling from generative autoencoders with Markov chains. *arXiv preprint arXiv:1610.09296*, 3:13, 2016.
- Arvanitidis, G., Hansen, L. K., and Hauberg, S. Latent space oddity: On the curvature of deep generative models. *arXiv:1710.11379 [stat]*, January 2018. URL <http://arxiv.org/abs/1710.11379>. arXiv: 1710.11379.
- Burda, Y., Grosse, R., and Salakhutdinov, R. Importance weighted autoencoders. *arXiv:1509.00519 [cs, stat]*, November 2016. URL <http://arxiv.org/abs/1509.00519>. arXiv: 1509.00519.
- Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., and Munafò, M. R. Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5): 365–376, 2013. Number: 5 Publisher: Nature Publishing Group.
- Calimeri, F., Marzullo, A., Stamile, C., and Terracina, G. Biomedical data augmentation using generative adversarial neural networks. In Lintas, A., Rovetta, S., Verschure, P. F., and Villa, A. E. (eds.), *Artificial Neural Networks and Machine Learning – ICANN 2017*, volume 10614, pp. 626–634. Springer International Publishing, Cham, 2017. ISBN 978-3-319-68611-0 978-3-319-68612-7. doi: 10.1007/978-3-319-68612-7_71. URL http://link.springer.com/10.1007/978-3-319-68612-7_71. Series Title: Lecture Notes in Computer Science.
- Caterini, A. L., Doucet, A., and Sejdinovic, D. Hamiltonian variational auto-encoder. In *Advances in Neural Information Processing Systems*, pp. 8167–8177, 2018.
- Chadebec, C., Mantoux, C., and Allasonnière, S. Geometry-aware Hamiltonian variational autoencoder. *arXiv:2010.11518 [cs, math, stat]*, October 2020. URL <http://arxiv.org/abs/2010.11518>. arXiv: 2010.11518.
- Chen, N., Klushyn, A., Kurle, R., Jiang, X., Bayer, J., and Smagt, P. Metrics for deep generative models. pp. 1540–1550. PMLR, 2018. ISBN 2640-3498.
- Cohen, G., Afshar, S., Tapson, J., and van Schaik, A. EMNIST: an extension of MNIST to handwritten letters. *arXiv:1702.05373 [cs]*, March 2017. URL <http://arxiv.org/abs/1702.05373>. arXiv: 1702.05373.
- Davidson, T. R., Falorsi, L., De Cao, N., Kipf, T., and Tomczak, J. M. Hyperspherical variational auto-encoders. *arXiv:1804.00891 [cs, stat]*, September 2018. URL <http://arxiv.org/abs/1804.00891>. arXiv: 1804.00891.
- Dilokthanakul, N., Mediano, P. A. M., Garnelo, M., Lee, M. C. H., Salimbeni, H., Arulkumaran, K., and Shanahan, M. Deep unsupervised clustering with Gaussian mixture variational autoencoders. *arXiv:1611.02648 [cs, stat]*, January 2017. URL <http://arxiv.org/abs/1611.02648>. arXiv: 1611.02648.
- Girolami, M. and Calderhead, B. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011. Number: 2 Publisher: Wiley Online Library.
- Girolami, M., Calderhead, B., and Chin, S. A. Riemannian manifold Hamiltonian Monte Carlo. *arXiv preprint arXiv:0907.1100*, 2009.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.
- Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. *Deep learning*, volume 1. MIT press Cambridge, 2016. Issue: 2.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-VAE: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2(5):6, 2017. Number: 5.
- Hoffman, M. D. and Johnson, M. J. Elbo surgery: yet another way to carve up the variational evidence lower bound. In *Workshop in Advances in Approximate Bayesian Inference, NIPS*, volume 1, pp. 2, 2016.
- Hsu, W.-N., Zhang, Y., and Glass, J. Unsupervised domain adaptation for robust speech recognition via variational autoencoder-based data augmentation. *arXiv:1707.06265 [cs]*, September 2017. URL <http://arxiv.org/abs/1707.06265>. arXiv: 1707.06265.

- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2261–2269, Honolulu, HI, July 2017. IEEE. ISBN 978-1-5386-0457-1. doi: 10.1109/CVPR.2017.243. URL <https://ieeexplore.ieee.org/document/8099726/>.
- Kalatzis, D., Eklund, D., Arvanitidis, G., and Hauberg, S. Variational autoencoders with Riemannian Brownian motion priors. *arXiv:2002.05227 [cs, stat]*, August 2020. URL <http://arxiv.org/abs/2002.05227>. arXiv: 2002.05227.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. *arXiv:1312.6114 [cs, stat]*, May 2014. URL <http://arxiv.org/abs/1312.6114>. arXiv: 1312.6114.
- Lebanon, G. Metric learning for text documents. *IEEE Trans. Pattern Anal. Machine Intell.*, 28(4):497–508, April 2006. ISSN 0162-8828. doi: 10.1109/TPAMI.2006.77. URL <http://ieeexplore.ieee.org/document/1597108/>.
- LeCun, Y. The MNIST database of handwritten digits. 1998. URL <http://yann.lecun.com/exdb/mnist/>.
- Lim, S. K., Loo, Y., Tran, N.-T., Cheung, N.-M., Roig, G., and Elovici, Y. Doping: Generative data augmentation for unsupervised anomaly detection with gan. pp. 1122–1127. IEEE, 2018. ISBN 1-5386-9159-0.
- Liu, Y., Zhou, Y., Liu, X., Dong, F., Wang, C., and Wang, Z. Wasserstein gan-based small-sample augmentation for new-generation artificial intelligence: a case study of cancer-staging data in biology. *Engineering*, 5(1):156–163, 2019. ISSN 2095-8099. Publisher: Elsevier.
- Mallasto, A. and Feragen, A. Wrapped Gaussian process regression on Riemannian manifolds. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5580–5588, Salt Lake City, UT, June 2018. IEEE. ISBN 978-1-5386-6420-9. doi: 10.1109/CVPR.2018.00585. URL <https://ieeexplore.ieee.org/document/8578683/>.
- Marcus, D. S., Wang, T. H., Parker, J., Csernansky, J. G., Morris, J. C., and Buckner, R. L. Open access series of imaging studies (OASIS): Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *Journal of Cognitive Neuroscience*, 19(9):1498–1507, 2007. doi: 10.1162/jocn.2007.19.9.1498. URL <https://doi.org/10.1162/jocn.2007.19.9.1498>. Number: 9 eprint: <https://doi.org/10.1162/jocn.2007.19.9.1498>.
- Mathieu, E., Le Lan, C., Maddison, C. J., Tomioka, R., and Teh, Y. W. Continuous hierarchical representations with poincaré variational auto-encoders. pp. 12565–12576, 2019.
- Nalisnick, E., Hertel, L., and Smyth, P. Approximate inference for deep latent Gaussian mixtures. pp. 4, 2016.
- Neal, R. M. Hamiltonian importance sampling. In *talk presented at the Banff International Research Station (BIRS) workshop on Mathematical Issues in Molecular Dynamics*, 2005.
- Neal, R. M. and others. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11):2, 2011. Number: 11.
- Nishizaki, H. Data augmentation and feature extraction using variational autoencoder for acoustic modeling. In *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pp. 1222–1227, Kuala Lumpur, December 2017. IEEE. ISBN 978-1-5386-1542-3. doi: 10.1109/APSIPA.2017.8282225. URL <http://ieeexplore.ieee.org/document/8282225/>.
- Pennec, X. Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *J Math Imaging Vis.*, 25(1):127–154, July 2006. ISSN 0924-9907, 1573-7683. doi: 10.1007/s10851-006-6228-4. URL <http://link.springer.com/10.1007/s10851-006-6228-4>.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Salimans, T., Kingma, D., and Welling, M. Markov chain Monte Carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, pp. 1218–1226, 2015.
- Sandfort, V., Yan, K., Pickhardt, P. J., and Summers, R. M. Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks. *Sci Rep*, 9(1):16884, December 2019. ISSN 2045-2322. doi: 10.1038/s41598-019-52737-x. URL <http://www.nature.com/articles/s41598-019-52737-x>.
- Shorten, C. and Khoshgoftaar, T. M. A survey on image data augmentation for deep learning. *J Big Data*, 6(1):60, December 2019. ISSN 2196-1115. doi: 10.1186/s40537-019-0197-0. URL <https://doi.org/10.1186/s40537-019-0197-0>.

[//journalofbigdata.springeropen.com/articles/10.1186/s40537-019-0197-0](http://journalofbigdata.springeropen.com/articles/10.1186/s40537-019-0197-0).

Tomczak, J. and Welling, M. VAE with a vampprior. pp. 1214–1223. PMLR, 2018. ISBN 2640-3498.

Turner, B. O., Paul, E. J., Miller, M. B., and Barbey, A. K. Small sample sizes reduce the replicability of task-based fMRI studies. *Communications Biology*, 1(1):1–10, 2018. Number: 1 Publisher: Nature Publishing Group.

Wu, Z., Wang, S., Qian, Y., and Yu, K. Data augmentation using variational autoencoder for embedding based speaker verification. In *Interspeech 2019*, pp. 1163–1167. ISCA, September 2019. doi: 10.21437/Interspeech.2019-2248. URL http://www.isca-speech.org/archive/Interspeech_2019/abstracts/2248.html.

Xiao, H., Rasul, K., and Vollgraf, R. Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. August 2017. arXiv: cs.LG/1708.07747.

Zhang, C., Bütepage, J., Kjellström, H., and Mandt, S. Advances in variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):2008–2026, 2018. Number: 8 Publisher: IEEE.

Zhu, X., Liu, Y., Qin, Z., and Li, J. Data augmentation in emotion classification using generative adversarial networks. *arXiv:1711.00648 [cs]*, December 2017. URL <http://arxiv.org/abs/1711.00648>. arXiv: 1711.00648.

Zhuang, P., Schwing, A. G., and Koyejo, O. Fmri data augmentation via synthesis. pp. 1783–1787. IEEE, 2019. ISBN 1-5386-3641-7.

A. Proof of Proposition. 1

Proposition 1. *The Riemannian manifold (\mathbb{R}^d, g) is geodesically complete.*

We will show that given the manifold $\mathcal{M} = \mathbb{R}^d$ and the Riemannian metric whose local representation is given by Eq. (3) in the paper, any geodesic curve $\gamma :]a, b[\rightarrow \mathcal{M}$ is actually extendible to \mathbb{R} . We derive a proof inspired from (Louis, 2019).

Proof. Let suppose that we have a geodesic curve γ such that γ cannot be extended to \mathbb{R} . There exists $I =]a, b[$ such that I is the maximum definition domain of γ . We will show that with such an assumption we end up with a contradiction.

Since L_{ψ_i} are lower triangular matrices with positive diagonal coefficients we have $x^\top L_{\psi_i} L_{\psi_i}^\top x > 0, \forall x \in \mathbb{R}^d - \{0\}$. Let $t_0 \in]a, b[$, we recall that

$$\|\dot{\gamma}(t)\|_{\gamma(t)}^2 = \langle \dot{\gamma}(t) | \dot{\gamma}(t) \rangle_{\gamma(t)} = \dot{\gamma}(t)^\top \mathbf{G}(\gamma(t)) \dot{\gamma}(t).$$

Therefore we have for any $t \in]a, b[$,

$$\begin{aligned} \|\dot{\gamma}(t)\|_2^2 &\leq \|\dot{\gamma}(t)\|_{\gamma(t)}^2 + \\ &\frac{1}{\lambda} \sum_{i=1}^N \dot{\gamma}(t)^\top L_{\psi_i} L_{\psi_i}^\top \dot{\gamma}(t) \exp\left(-\frac{\|\gamma(t) - c_i\|_2^2}{T^2}\right) \\ &\leq \frac{1}{\lambda} \cdot \|\dot{\gamma}(t)\|_{\gamma(t)}^2 = \frac{1}{\lambda} \cdot \|\dot{\gamma}(t_0)\|_{\gamma(t_0)}^2, \end{aligned}$$

where the last equality comes for the constant speed of geodesic curves. Hence we have:

$$\|\gamma(t) - \gamma(t_0)\|_2 \leq \frac{\|\dot{\gamma}(t_0)\|_{\gamma(t_0)}}{\sqrt{\lambda}} \cdot |t - t_0|.$$

This means that for any $t \in]a, b[$ the geodesic curve γ remains within a compact set.

We now show that the curve can actually be extended. Let us define the sequence $t_n \xrightarrow{n \rightarrow \infty} b$. Since the geodesic curves have a constant speed the set $I = \{(t_n, \dot{\gamma}(t_n))\}_{n \in \mathbb{N}}$ is compact. Moreover, using Cauchy-Lipchitz theorem, we can find $\varepsilon > 0$ such that for any $n \in \mathbb{N}$, the geodesic γ can be defined on $]t_n - \varepsilon, t_n + \varepsilon[$. Since t_n can be as close to b as desired, there exists $N \in \mathbb{N}$ such that $\forall n \geq N$ we have $t_n \geq b - \frac{\varepsilon}{2}$. This means that the curve definition domain can be extended to $]a, b + \frac{\varepsilon}{2}[$ which concludes the proof. \square

B. Discussion of Remark. 1

Remark 1. *If Σ has small enough eigenvalues then Alg. 1 of Sec. 4.3 samples from the following distribution on \mathbb{R}^d*

$$p(z) = \frac{\rho_S(z) \sqrt{\mathbf{G}^{-1}(z)}}{\int_{\mathbb{R}^d} \rho_S(z) \sqrt{\mathbf{G}^{-1}(z)} dz}, \quad (1)$$

where $\rho_S(z) = 1$ if $z \in S$, 0 otherwise, and S is the support of the distribution taken as a compact set so that the integral is well defined.

If Σ has small enough eigenvalues, it means that the initial velocity $v \sim \mathcal{N}(0, \Sigma)$ will have a low magnitude with high probability. In such a case, we can show with some approximation that the ratio α in the Riemannian random walk presented in Sec. 4.3 is a Hasting-Metropolis ratio with target density given by Eq. (1). We recall that the classic Hasting-Metropolis ratio writes

$$\alpha(x, y) = \frac{\pi(y)}{\pi(x)} \cdot \frac{q(x, y)}{q(y, x)},$$

where π is the target distribution and q a proposal distribution. In the case of small magnitude velocities, one may show that q is symmetric that is

$$q(x, y) = q(y, x).$$

In our setting, a proposal \tilde{z} is made by computing the geodesic γ starting at $\gamma(0) = z$ with initial velocity $\dot{\gamma}(0) = v$ where $v \sim \mathcal{N}(0, \Sigma)$ and evaluating it at time 1. First, we remark that γ is well defined since the Riemannian manifold $\mathcal{M} = (\mathbb{R}, g)$ is *geodesically complete* and γ is unique. Moreover, we have that $\overleftarrow{\gamma}$ is the unique geodesic with initial position $\overleftarrow{\gamma}(0) = \tilde{z} = \gamma(1)$ and initial velocity $\overleftarrow{\dot{\gamma}}(0) = -\dot{\gamma}(1)$ and we have

$$\overleftarrow{\dot{\gamma}}(t) = \dot{\gamma}(1 - t), \quad \forall t \in [0, 1].$$

In the case of small enough initial velocity, a Taylor expansion of the exponential may be performed next to $0 \in T_z \mathcal{M}$ and consists in approximating geodesic curves with straight lines. That it, for $t \in [0, 1]$

$$\text{Exp}_z(vt) \approx z + vt,$$

where $v = \tilde{z} - z$. In such a case we have

$$\dot{\gamma}(t) = v = \dot{\gamma}(0) = \dot{\gamma}(1) = -\overleftarrow{\dot{\gamma}}(0).$$

Moreover, we have on the one hand

$$q(\tilde{z}, z) = p(\tilde{z}|z) = p(\text{Exp}_z(v)|z) \simeq \mathcal{N}(z, \Sigma).$$

On the other hand

$$q(z, \tilde{z}) = p(z|\tilde{z}) = p(\text{Exp}_{\tilde{z}}(-v)|\tilde{z}) \simeq \mathcal{N}(\tilde{z}, \Sigma)$$

Therefore,

$$q(\tilde{z}, z) = q(z, \tilde{z}).$$

Finally, the ratio α in the Riemannian random walk may be seen as a Hasting-Metropolis ratio where the target density is given by Eq. (1) and so the algorithm samples from such a distribution.

C. Computing the Exponential map

To compute the exponential map at any given point $p \in \mathcal{M}$ and for any tangent vector $v \in T_p\mathcal{M}$ we rely on the Hamiltonian definition of geodesic curves. First, for any given $v \in T_p\mathcal{M}$, the linear form:

$$q_v : \begin{cases} T_p\mathcal{M} & \rightarrow \mathbb{R} \\ u & \rightarrow g_p(u, v) \end{cases},$$

is called a moment and is a representation of v in the dual space. In short, we may write $q_v(u) = u^\top \mathbf{G}v$. Then, the definition of the Hamiltonian follows

$$H(p, q) = \frac{1}{2} g_p^*(q, q),$$

where g_p^* is the dual metric whose local representation is given by $\mathbf{G}^{-1}(p)$, the inverse of the metric tensor. Finally, all along geodesic curves the following equations hold

$$\frac{\partial p}{\partial t} = \frac{\partial H}{\partial q}, \quad \frac{\partial q}{\partial t} = -\frac{\partial H}{\partial p}. \quad (2)$$

Such a system of differential equations may be integrated pretty straight forwardly using simple numerical schemes such as the second order Runge Kutta integration method and Alg. 1 as in (Louis, 2019). An implementation using pytorch (Paszke et al., 2017) is made available in the supplementary materials. Noteworthy is the fact that such an algorithm only involves one metric tensor inversion at initialisation to recover the initial moment from the initial velocity. Moreover, it involves closed form operations since the inverse metric tensor \mathbf{G}^{-1} is known and so the gradients in Eq. (2) can be easily computed.

Algorithm 1 Computing the Exponential map

Input: $z_0 \in \mathcal{M}$, $v \in T_{z_0}\mathcal{M}$ and T
 $q \leftarrow \mathbf{G} \cdot v$
 $dt \leftarrow \frac{1}{T}$
for $t = 1 \rightarrow T$ **do**
 $p_{t+\frac{1}{2}} \leftarrow p_t + \frac{1}{2} \cdot dt \cdot \nabla_q H(p_t, q_t)$
 $q_{t+\frac{1}{2}} \leftarrow p_t - \frac{1}{2} \cdot dt \cdot \nabla_p H(p_t, q_t)$
 $p_{t+1} \leftarrow p_t + dt \cdot \nabla_q H(p_{t+\frac{1}{2}}, q_{t+\frac{1}{2}})$
 $q_{t+1} \leftarrow p_t - dt \cdot \nabla_p H(p_{t+\frac{1}{2}}, q_{t+\frac{1}{2}})$
end for
Return p_T

D. Paper Parameters Setting

Table. 1 summarises the main hyper-parameters we use to perform the experiments presented in the paper while Table. 2 shows the neural networks architectures employed. As to training parameters, we use a Adam optimiser (Kingma & Ba, 2014) and the learning rate is set to 10^{-3} for all data sets

except for OASIS where it is decreased to 10^{-4} and is the same for each model. As to the models used as benchmark for data augmentation, the DenseNet implementation we use is the one in (Amos, 2020) with a *growth rate* equals to 10, *depth* of 20 and 0.5 *reduction* and is trained with a learning rate of 10^{-3} . For OASIS, the MLP has 400 hidden units and relu activation function.

Table 1. RHVAE parameters for each data set.

DATA SETS	PARAMETERS					
	d^*	n_{lf}	ϵ_{lf}	T	λ	$\sqrt{\beta_0}$
SYNTHETIC	2	3	10^{-2}	0.8	10^{-3}	0.3
reduced MNIST	2	3	10^{-2}	0.8	10^{-3}	0.3
reduced FASHION	2	3	10^{-2}	0.8	10^{-3}	0.3
OASIS	2	3	10^{-3}	0.8	10^{-2}	0.3

* LATENT SPACE DIMENSION (SAME FOR VAE AND VAMP-VAE)

Table 2. Neural networks architectures of the VAE, VAMP-VAE and RHVAE for each data set. The *encoder* and *decoder* are the same for all models.

SYNTHETIC, MNIST & FASHION			
NET	LAYER 1	LAYER 2	LAYER 3
μ_ϕ^* Σ_ϕ^*	(D , 400, RELU)	(400, d , LIN.) (400, d , LIN.)	-
π_θ^*	(d , 400, RELU)	(400, D , SIG.)	-
L_ψ (DIAG.) L_ψ (LOW.)	(D , 400, RELU)	(400, d , LIN.) (400, $\frac{d(d-1)}{2}$, LIN.)	-
OASIS			
μ_ϕ^* Σ_ϕ^*	(D , 1K, RELU)	(1K, 400, RELU)	(400, d , LIN.) (400, d , LIN.)
π_θ^*	(d , 400, RELU)	(400, 1K, RELU)	(1K, D , SIG.)
L_ψ (DIAG.) L_ψ (LOW.)	(D , 400, RELU)	(400, d , LIN.) (400, $\frac{d(d-1)}{2}$, LIN.)	-

* SAME FOR VAE MODELS

E. Comparison with Other Generation Methods

In this section we compare the visual quality of generated samples using *prior-based* methods (\mathcal{N} -VAE, VAMP-VAE and RHVAE) and the proposed methods (RHVAE + Riemannian random walk or RHVAE + metric sampling). We compare the models on 1) the synthetic data set composed by 180 binary images presenting circles and rings with different diameters and thicknesses; 2) the *reduced* MNIST data set composed by 120 samples of 3 different classes and 3) the *reduced* FashionMNIST data set composed again by 120 samples from 3 distinct classes. The models are trained with the same neural network architectures, batch size and learning rate. An early stopping strategy is adopted and consists in stopping training if the ELBO does not improve for 50 epochs.

As discussed in the paper, changing the prior may indeed improve the model generation capacity. For instance samples from the VAE with the VAMP prior (3rd row of Fig. 1) are closer to the training data (1st row of Fig. 1) than with the gaussian prior (2nd and 4th row). The model is for instance able to generate circles when trained with the synthetic data while models using a standard normal prior are not. Nonetheless, a non negligible part of the generated samples are degraded (see saturated images for the *reduced* MNIST data for instance). This aspect is mitigated with the proposed generation methods which generate more diverse and sharper samples.

F. Comprehensive Metrics Classification of Sec. 5.2

In addition to the global and balanced accuracy metrics provided in the paper, we record the precision, recall and f1-score for each class of the OASIS database (*i.e.* Alzheimer disease (AD) and Cognitively Normal (CN) patients) in Table. 3. Again, the best results are observed with the *synthetic only* data sets only composed by samples generated with our method (*i.e.* metric sampling). For instance, f1-score jumps from 66.89 with the *raw data* to 78.05 with 1000 synthetic samples per class using our method for the AD diagnosis. This represents a noticeable improvement considering the size of the original data set (416 patients) and the complexity of the data (208x176 RMI). Noteworthy is the fact that in the majority of cases we are also able to reduce the model uncertainty with the proposed augmentation method. We recall that the classifier is trained on 5 independent runs and that a hyper-parameter search is conducted on the learning rate which appears to have a significant impact on the results. We keep the learning rate $\in \{10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$ achieving the best balanced accuracy.

References

- Amos, B. bamos/densenet.pytorch, December 2020. URL <https://github.com/bamos/densenet.pytorch>. original-date: 2017-02-09T15:33:23Z.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Louis, M. *Computational and statistical methods for trajectory analysis in a Riemannian geometry setting*. PhD Thesis, Sorbonnes universités, 2019.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch. 2017.

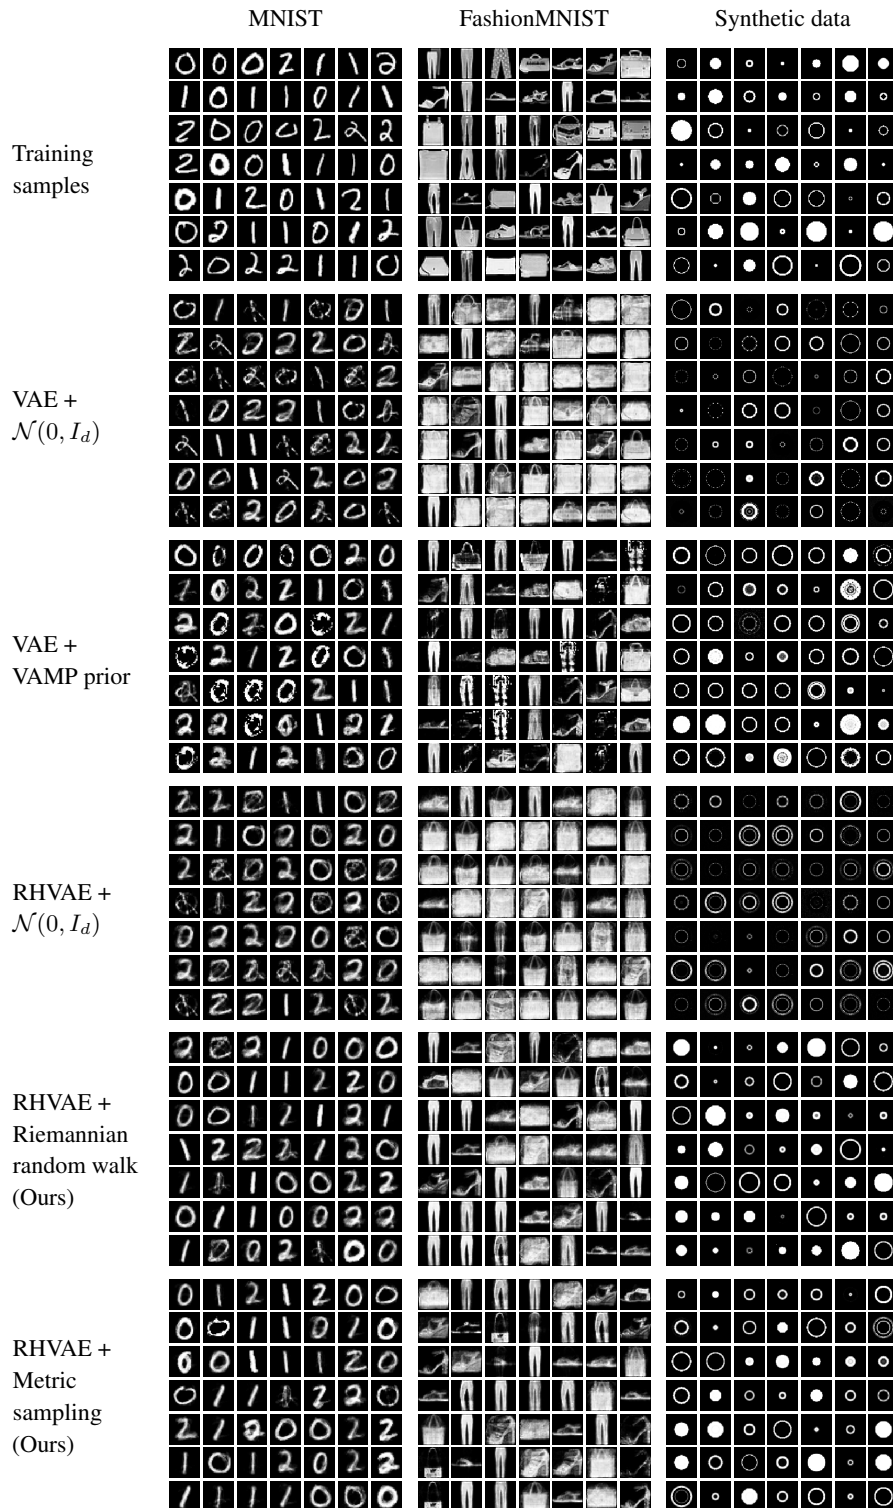


Figure 1. Comparison of 5 sampling methods on the *reduced* MNIST, *reduced* FashionMNIST and the synthetic data sets. From top to bottom: 1) samples extracted from the training set; 2) samples generated with a Vanilla VAE and using the prior; 3) from the VAMP prior VAE ; 4) from a RHVAE and the *prior-based* generation scheme; 5) from a RHVAE and using the proposed Riemannian random walk and 6) using metric sampling. All the models are trained with the same encoder and decoder networks and identical latent space dimension. An early stopping strategy is adopted and consists in stopping training if the ELBO does not improve for 50 epochs.

Table 3. Data augmentation on real data from the OASIS database with a MLP model as Benchmark using either geometric transformations, a VAE and the classic generation process or metric sampling (ours). The model is trained on 5 independent runs and we report the main classification metrics along with the associated standard deviation. We adopt an early stopping strategy and stop training if the validation loss does not improve for 20 epochs.

DATA SETS	LABEL	PRECISION	RECALL	F1-SCORE	GLOBAL ACCURACY	BALANCED ACCURACY
RAW DATA	CN	92.61 (2.1)	86.97 (1.2)	89.69 (1.4)	84.29 (2.2)	80.71 (4.0)
	AD	60.82 (3.8)	74.44 (7.5)	66.89 (5.2)		
REAL + SYNTHETIC						
AUGMENTED X5	CN	94.71 (2.2)	85.45 (3.0)	89.79 (1.6)	84.76 (2.3)	83.84 (3.8)
	AD	60.91 (4.5)	82.22 (8.2)	69.76 (4.7)		
AUGMENTED X10	CN	90.22 (1.9)	89.09 (0.6)	89.64 (1.1)	83.81 (1.8)	76.77 (3.8)
	AD	61.53 (3.1)	64.44 (7.5)	62.87 (5.1)		
AUGMENTED X15	CN	90.37 (2.1)	86.97 (4.1)	88.54 (1.6)	82.38 (2.0)	76.26 (3.0)
	AD	58.74 (5.9)	65.56 (8.9)	61.38 (3.3)		
VAE						
VAE - 200*	CN	92.74 (1.7)	83.64 (4.8)	87.83 (2.0)	81.90 (2.3)	79.59 (1.6)
	AD	56.64 (5.0)	75.56 (7.5)	64.20 (1.5)		
VAE - 500*	CN	91.58 (0.5)	92.12 (2.4)	91.83 (1.0)	87.14 (1.4)	80.50 (0.5)
	AD	71.07 (5.2)	68.89 (2.7)	69.74 (1.5)		
VAE - 1000*	CN	93.33 (0.0)	84.85 (0.0)	88.89 (0.0)	83.33 (0.0)	81.31 (0.0)
	AD	58.33 (0.0)	77.78 (0.0)	66.67 (0.0)		
VAE - 2000*	CN	93.24 (0.0)	83.64 (0.6)	88.18 (0.4)	82.38 (0.5)	80.71 (0.3)
	AD	56.47 (0.9)	77.78 (0.0)	65.43 (0.6)		
OURS						
OURS - 200*	CN	96.28 (1.3)	83.94 (8.7)	89.40 (4.7)	84.76 (6.0)	85.86 (3.0)
	AD	62.79 (12.0)	87.78 (5.4)	72.15 (7.1)		
OURS - 500*	CN	95.71 (1.0)	86.67 (4.4)	90.88 (2.0)	86.43 (2.7)	86.11 (1.0)
	AD	64.67 (6.9)	85.56 (4.4)	73.23 (3.0)		
OURS - 1000*	CN	95.71 (1.7)	86.36 (3.7)	90.73 (1.8)	86.19 (2.3)	85.96 (2.6)
	AD	63.74 (5.2)	85.56 (6.7)	72.71 (3.4)		
OURS - 2000*	CN	94.34 (1.2)	90.61 (0.6)	92.43 (0.2)	88.33 (0.5)	85.30 (1.9)
	AD	69.91 (0.2)	80.00 (4.4)	74.56 (1.7)		
SYNTHETIC ONLY						
VAE - 200*	CN	96.66 (1.0)	70.00 (0.6)	81.20 (0.7)	74.52 (0.9)	80.55 (1.6)
	AD	45.30 (1.1)	91.11 (2.7)	60.51 (1.6)		
VAE - 500*	CN	91.58 (1.9)	86.67 (8.4)	88.75 (3.8)	83.10 (4.8)	78.33 (1.9)
	AD	62.78 (11.8)	70.00 (9.7)	64.51 (4.3)		
VAE - 1000*	CN	95.92 (0.0)	71.21 (0.0)	81.74 (0.0)	75.00 (0.0)	80.05 (0.0)
	AD	45.71 (0.0)	88.89 (0.0)	60.38 (0.0)		
VAE - 2000*	CN	95.65 (0.7)	73.03 (0.6)	82.82 (0.1)	76.19 (0.0)	80.41 (0.8)
	AD	47.02 (0.1)	87.78 (2.2)	61.23 (0.6)		
OURS						
OURS - 200*	CN	95.22 (0.0)	90.61 (0.6)	92.86 (0.3)	89.05 (0.5)	86.97 (0.3)
	AD	70.78 (1.3)	83.33 (0.0)	76.54 (0.8)		
OURS - 500*	CN	96.41 (0.6)	89.39 (0.0)	92.77 (0.3)	89.05 (0.5)	88.58 (1.1)
	AD	69.29 (0.5)	87.78 (2.2)	77.44 (1.2)		
OURS - 1000*	CN	96.72 (0.0)	89.39 (0.0)	92.91 (0.0)	89.29 (0.0)	89.14 (0.0)
	AD	69.57 (0.0)	88.89 (0.0)	78.05 (0.0)		
OURS - 2000*	CN	96.41 (0.6)	89.39 (0.0)	92.77 (0.3)	89.05 (0.5)	88.58 (1.1)
	AD	69.29 (0.5)	87.78 (2.2)	77.44 (1.2)		

* NUMBER OF GENERATED SAMPLES