

---

# Learning the Base Distribution in Implicit Generative Models

---

## Abstract

Popular density estimation methods such as Generative Adversarial Networks (GANs), and Variational Autoencoders (VAE) enforce the latent representation to follow simple distributions such as isotropic Gaussian. In this paper, we claim that learning a complicated distribution over the latent space of an auto-encoder enables accurate modeling over complicated data distributions. We propose a two stage optimization procedure which approximately maximizes the implicit density model. We experimentally verify that our method can generate realistic samples on two image datasets (MNIST, CELEB-A) and one audio dataset (spoken-digit dataset)

## 1 Introduction

Generative model learning is the task where the goal is to learn a model to generate artificial samples which follow the underlying probability density function of a given dataset. When the dataset comprises of scalars, or of low dimensional (2-3 dimensions) vectors and follow a uni-modal distribution, one can use a simple density model such as multivariate Gaussian, and fit the model on the data using maximum likelihood. To learn the distribution of more complicated data such as natural images, or audio simple densities do not have the required expressive power, because of the aforementioned high dimensionality, multi-modality and structured nature of the data.

There exists several main generative model learning methods in the machine learning literature. One way of approaching the problem is to use a linear latent variable model such as mixture models [1], a latent factor model such as probabilistic PCA [2], Hidden Markov model [3], or linear dynamical systems [4]. These mod-

els can successfully capture the multi-modality, or low rank nature of the datasets, however they rely on linear and tractable forward mappings, and therefore lack the expressive power of modern neural network models.

More recently, the mainstream approaches for learning a generative model for complicated datasets have been centered around models that combine latent variable modeling with non-linear neural network mappings. One prominent example of such approaches is Variational Autoencoders (VAEs) [5]. VAEs consider a latent variable model where the latent representation is mapped to the observation space via a complicated neural network. The variational expectation maximization algorithm in [5] maximize a variational lower bound on the maximum likelihood objective. The prior distribution is typically chosen as a simple distribution such that the KL-divergence term in the lower bound is tractable. In this paper we argue that using a simple prior distribution is detrimental to the overall quality of the learned generative model.

Another very popular method that also uses a restricted latent representation is Generative Adversarial Networks (GANs) [6]. The main conceptual differences of GANs from typical latent variable models (including VAEs) is that GANs are an implicit generative model learning methodology [7], where the model distribution is defined without specifying an output density. More importantly, unlike LVMs GANs do not maximize the standard maximum likelihood objective. Instead, GANs approximate the underlying dataset density via an additional discriminator network. Although an appealing idea, GANs are incredibly hard to train (as evidenced by the sheer number of GAN training papers in the last few years), and suffer from the predictable mode collapse problem (We delve more into this in the main text).

In this paper, we propose an implicit generative model learning method which maximizes the maximum likelihood training objective. Unlike GANs, the method

does not rely on auxiliary networks such as discriminator or critic networks. For training, we propose a simple two stage training method, which maximizes a maximum likelihood training objective, and therefore does not suffer from the mode collapse problem that GANs are well known for.

## 2 Generative Model Learning

The purpose of this section is to set the notation and the required concepts before we formally introduce our algorithm. As we discussed in the introduction, the goal in generative model learning is to approximate the underlying data density  $p_{\text{data}}(x)$  with the density that our model implies, which we denote by  $p_{\text{model}}(x|\theta)$ , where  $\theta$  denotes the model parameters. The maximum likelihood training minimizes the Kullback-Leibler (KL) divergence between the data density and model density:

$$\begin{aligned} \min_{\theta} \text{KL}(p_{\text{data}}(x) \| p_{\text{model}}(x|\theta)) \\ &= \int p_{\text{data}}(x) \log \frac{p_{\text{data}}(x)}{p_{\text{model}}(x|\theta)} dx \\ &\propto - \int p_{\text{data}}(x) \log p_{\text{model}}(x|\theta) dx \\ &\propto - \sum_n \log p_{\text{model}}(x_n|\theta) \end{aligned} \quad (1)$$

where we recognize Equation (1) as the negative of the maximum likelihood objective. Note that  $x \in \mathbb{R}^L$  denotes the variable we use to denote the observation space, and we use the subscripted version  $x_n$  to denote the data item with index  $n$ .

It is usually not possible to get an easy to compute (tractable) likelihood function  $p_{\text{model}}(x|\theta)$  unless we work with very simple models. In LVMs, Jensen's inequality is used to compute a lower bound to the maximum likelihood objective:

$$\begin{aligned} \log p_{\text{model}}(x|\theta) &= \log \int p_{\text{forward}}(x|h, \theta) p(h) dh, \\ &= \log \int \frac{p_{\text{forward}}(x|h, \theta) p(h)}{q(h)} q(h) dh, \\ &\geq \mathbb{E}_{q(h)} [\log p_{\text{forward}}(x|h, \theta)] - \text{KL}(q(h) \| p(h)), \end{aligned} \quad (2)$$

where Equation (2) is known as the variational lower bound, or ELBO [1],  $h \in \mathbb{R}^K$  denotes the latent variable, and  $q(h)$  denotes the variational distribution over the latent variable. In linear LVMs with tree structured latent variables (e.g. mixture models, HMMs), we can use the posterior  $p(h|\theta)$  as the variational distribution, because the posterior makes this bound tight.

In the general situation where the forward mapping is defined via a non-linear mapping, such that

$p_{\text{forward}}(x|h, \theta) = p_{\text{out}}(x; f_{\theta}(h))$ , where  $f_{\theta}(h) : \mathbb{R}^K \rightarrow \mathbb{R}^L$  is the nonlinear deterministic mapping, and  $p_{\text{out}}(\cdot)$  is the employed noise model, computing the posterior distribution is not analytically possible. VAEs therefore use a neural network mapping for the variational distribution such that  $q_{\phi}(h) = \mathcal{N}(x; \mu_{\phi}(x), \sigma_{\phi}^2(x)I)$ , where  $\mathcal{N}(\cdot)$  denotes the Normal distribution and the neural network mappings  $\mu_{\phi}(x), \sigma_{\phi}^2(x) : \mathbb{R}^L \rightarrow \mathbb{R}^K$  parametrize the variational distribution.

Although the likelihood computation in VAEs is intractable and require the variational EM algorithm described above, we argue in this paper that the main failure mode of VAEs is caused by the simplistic prior choices for  $p(h)$ . We demonstrate this in the experiments section.

Another popular way to learn generative models is via GANs. GANs are implicit generative models, therefore they do not employ an output distribution  $p_{\text{out}}$ . Namely, the data generation mechanism is defined as follows:

$$h \sim p^0(h), \quad x = f_{\theta}(h), \quad (3)$$

where we call  $p^0(h)$  the *base distribution*, which is typically chosen as a simplistic distribution such as an isotropic Gaussian distribution, and  $f_{\theta}(h)$  is a deterministic forward mapping similar to what we have denoted for VAEs above. GANs therefore do not employ an output distribution  $p_{\text{out}}(\cdot)$ , but rather define  $p_{\text{model}}(x|\theta)$  via a deterministic transformation of the base distribution  $p^0(h)$ .

In this paper, we also argue that one of the reasons why GANs might underperform is because of the simplistic base distribution choice. In addition to this, GANs also complicate the learning process (model parameter optimization) by introducing a discriminator network. GANs in their original formulation [2], approximate the ratio between the data density and the model density [4]:

$$\begin{aligned} \mathcal{L}(\theta, \xi) &= \sum_n \log D_{\xi}(x_n) + \sum_{n'} \log 1 - D_{\xi}(x'_{n'}) \\ &\rightarrow \sum_n \log \frac{p_{\text{data}}(x_n)}{p_{\text{model}}(x_n|\theta) + p_{\text{data}}(x_n)} \\ &\quad + \sum_{n'} \log \frac{p_{\text{model}}(x'_{n'}|\theta)}{p_{\text{model}}(x'_{n'}|\theta) + p_{\text{data}}(x'_{n'})} \end{aligned} \quad (4)$$

where  $x_n$  denotes the training instances, and  $x'_{n'}$  denotes samples generated from the model. The convergence to the second line (which can be recognized as the Monte Carlo estimate for the Jensen-Shannon divergence) can be easily seen by maximizing the objective  $\mathcal{L}(\theta, \xi)$  with respect to the discriminator parameters  $\xi$  [2]. One other big conceptual problem with GANs is that when we minimize this divergence with respect to  $\theta$ , it is well known

Jensen-Shannon divergence collapses down to a subset of the modes of the data distribution  $p_{\text{data}}(x)$ .

There exists several other variants of GANs where instead of using the Jensen-Shannon divergence [all the fdvergence stuff], one can consider Wasserstein distance [1], optimal transport objectives [ ], or other tricks to make the optimization easier [schwing paper], or alleviate the mode collapse problem [gan training trick paper]. But there has not been enough focus on achieving a more expressive distribution via learning the base distribution.

### 3 Learning in Implicit Generative Models

As we mentioned before GANs are only one option for learning implicit generative models. According to [?], there are three other broad categories which are ratio matching [ ], divergence matching [ ] and moment matching [ ]. The first two methods are basically variations on the GAN score function, and do not consider the effect of base distribution. Moment matching is a simple method where the observable moments of the model density are matched with the moments computed from data. In our experience this method's performance highly depends on the quality of the choice of the moments. In this paper, we propose a novel maximum likelihood based implicit generative model learning method, which focuses on learning expressive base distributions.

We know from probability theory that in an implicit generative model, the output probability density is related to the base distribution via the cumulative density function:

$$p_{\text{model}}(x|\theta, \phi) = \frac{\partial}{\partial x} \int_{\{x: f_{\theta}(h) \leq x\}} p_{\phi}^0(h) dh, \quad (5)$$

where note that the base distribution is parametrized by  $\phi$ . The integral in Equation (5) is not tractable in general, however if we have an invertible mapping  $f_{\theta}(h)$ , we can obtain an analytical expression for the density function of the model using the following formula:

$$p_{\text{model}}(x|\theta, \phi) = p_{\phi}^0(f_{\theta}^{-1}(x)) |\det J_{\theta}(x)|, \quad (6)$$

where  $J_{\theta}(x) := \frac{\partial f_{\theta}^{-1}(x)}{\partial x}$ . It is possible to construct exactly invertible mappings using typical neural network mappings such as matrix multiplications and convolutions. In this paper we explore the cases where the mapping  $f_{\theta}(x)$  is invertible in section ???. We also look at the cases where  $f_{\theta}(h)$  is approximately invertible, and propose an algorithm which approximately maximizes the maximum likelihood objective.

#### 3.1 Maximum Likelihood for Implicit Generative Models

If we work with invertible forward mappings, the optimization problem for maximum likelihood in an implicit generative model is the following:

$$\begin{aligned} & \max_{\theta, \phi} \sum_n \log p_{\text{model}}(x_n|\theta, \phi), \\ & = \max_{\theta, \phi} \sum_n \log p_{\phi}^0(f_{\theta}^{-1}(x_n)) + \log \det |J_{\theta}(x_n)|, \quad (7) \end{aligned}$$

where the goal can be interpreted as maximizing the likelihood of the mappings  $f_{\theta}^{-1}(x)$  in the base distribution space. The determinant of the Jacobian term ensures that the distribution properly normalized. If we think about this objective from a sampling perspective, in order to the generate plausible samples, the maximum likelihood objective tries to match the samples from base distribution with the observations mapped to the base distribution space  $f_{\theta}^{-1}(x)$ .

Note that in GANs, only the forward mapping parameters  $\theta$  is optimized, and the base distribution is fixed to be simple unimodal distribution. Optimizing both the forward mapping and a multi-modal base distribution constitutes the main idea in our paper. We argue that mapping a multimodal dataset onto a unimodal base distribution is harder to achieve than fitting a multimodal distribution on  $f_{\theta}^{-1}(x)$ . We demonstrate this in Figure 1. Using an invertible linear mapping  $f_{\theta}(h) = Wh$ , where  $h \in \mathbb{R}^2$ , and  $W \in \mathbb{R}^{2 \times 2}$ , we show that on a two dimensional mixture of Gaussians example that, if we do maximum likelihood on the objective in Equation (8), we fail to map the observations to the samples drawn from a fixed isotropic base distribution. However, as shown in Figure (b) if we set the base distribution as a flexible distribution such as mixture of Gaussians, and learn its parameters  $\phi$ , we are able to learn a much more accurate distribution. We also show that if we train the same mapping using the standard GAN formulation, we get the mode collapse behavior, where only one of the Gaussians is captured in the learned distribution.

We acknowledge that in the cases where the forward mapping has the same dimensionality in the domain and range spaces (such as the example in Figure 1), learning an implicit generative model by maximizing Equation (8) is pointless, because we could have very well just fitted a mixture model on the data. For this reason, in the next section we propose the two stage learning algorithm which learns an expressive base distribution on a low dimensional space.

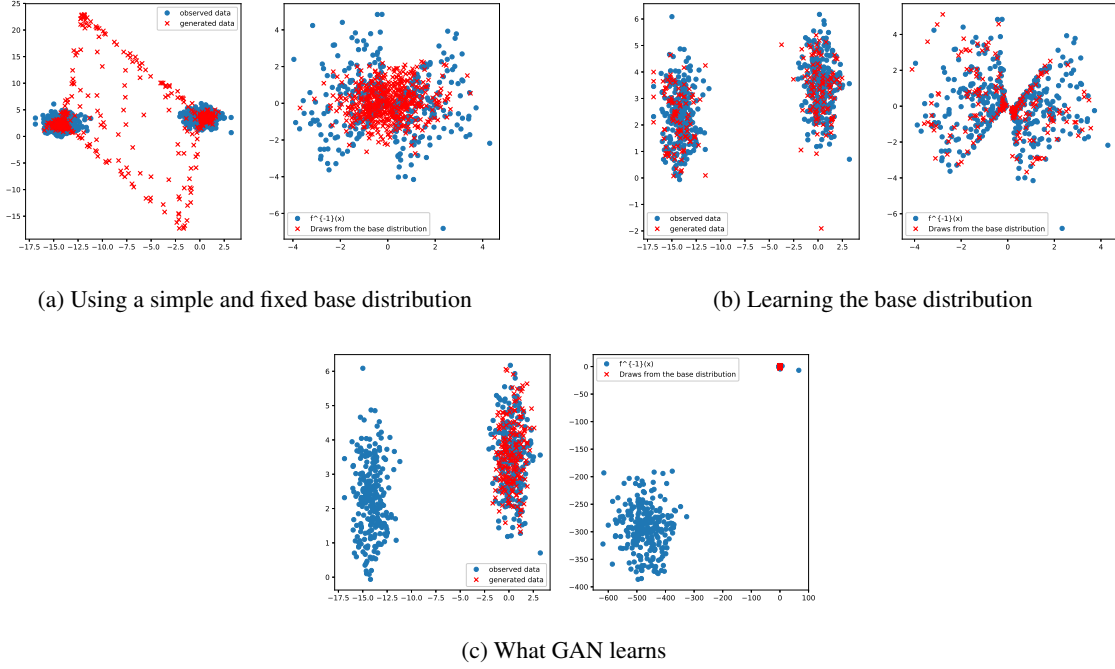


Figure 1: We demonstrate the differences between our proposed method and two methods that use a simple base distributions on a two dimensional mixture of two Gaussians. (Maximizing (8) and GAN with a fixed isotropic Gaussian as base distribution). In figure (a) we maximize the implicit model likelihood defined in (8) with respect to forward mapping parameters  $\theta$ . In figure (b) we fit a mixture of Gaussians for the same forward mapping as figure (a). In figure (c) we see what GAN learns for the same dataset. In Each figure, left plot shows the generated data overlaid on the observed dataset, and left plot show the samples from the base distribution overlaid on the observations mapped to the observation space ( $f_\theta^{-1}(x)$ ).

### 3.2 The Two Stage Algorithm

In practice, we typically would like to have base distribution defined on a space which has lower dimensionality than the observation space. If this is the case, then it is impossible to have an exactly invertible mapping  $f_\theta(h)$ . It is however possible to have an approximately invertible forward mapping. This idea gives the hint for a very simple two stage maximum likelihood algorithm: We first fit an auto-encoder such that the error  $\sum_n \|f_\theta(f_\psi^{\text{enc}}(x_n)) - x_n\|$  is minimized. Once the we are done with optimizing the autoencoder, we simply fit a base distribution on the embeddings  $f_\psi^{\text{enc}}(x)$ . The formal algorithm is specified in Algorithm 1.

To see that this is a maximum likelihood algorithm, let us reconsider the likelihood function of the implicit generative model with the autoencoder:

$$= \max_{\phi} \sum_n \log p_{\phi}^0(f_{\psi}^{\text{enc}}(x_n)) + \log \det |\tilde{J}_{\theta, \phi}(x_n)|, \quad (8)$$

where we easily see that the base distribution parameters are independent of the Jacobian term. Assuming that that

---

**Algorithm 1** The two stage implicit generative model learning algorithm

---

-Train the auto-encoder parameters  $\theta, \psi$  such that:

$$\min_{\theta, \psi} \sum_n \|f_{\theta}(f_{\psi}^{\text{enc}}(x_n)) - x_n\|$$

-Fit the base distribution on the latent space such that:

$$\max_{\phi} \sum_n \log p_{\phi}^0(f_{\psi}^{\text{enc}}(x_n))$$


---

the autoencoder learns a mapping close to the identity, we conclude that maximizing with respect to the base distribution parameters maximizes the model likelihood.

Note that since the optimization for the forward mapping parameters  $\theta$ , and the base distribution is decoupled, it is easy to fit a multi-modal distribution for the base distribution on the embeddings  $f_{\psi}^{\text{enc}}(x)$ . One natural choice is to use a mixture distribution. We demonstrate this on handwritten zero and one digits from the MNIST dataset

[ ] in Figure 2. We choose the dimensionality of the latent space  $K = 2$  to be able to visualize the base distribution space. We use a three component Gaussian mixture model for this example.

### 3.3 Learning Generative Models for Sequential Data

The framework we propose also offers the flexibility to learn distributions over sequences by simply learning a sequential distribution such as HMM on the latent representations. The likelihood of a sequence is expressed as follows:

$$p_{\text{model}}(x_{1:T}|\psi, \phi) = \prod_{t=1}^T p_{\phi}^0(f_{\psi}^{\text{enc}}(x_t)|f_{\psi}^{\text{enc}}(x_{1:T-1})) \times |\det J_{\theta}(x)|,$$

where a sequence is denoted as  $x_{1:T} := \{x_1, x_2, \dots, x_T\}$ . and thus  $f_{\psi}^{\text{enc}}(x_{1:T-1}) = \{f_{\psi}^{\text{enc}}(x_1), f_{\psi}^{\text{enc}}(x_2), \dots, f_{\psi}^{\text{enc}}(x_{T-1})\}$ . According to this density model, the observations  $x_{1:T}$  are mapped to latent space independent from each other. This suggests that we can closely follow the two stage algorithm defined in Algorithm 1: Same as before we first fit the autoencoder, and obtain the latent representations. In the second stage, instead of fitting an exchangeable model such as a mixture model, we fit a base distribution which models the temporal structure of the latent space. Potential options for such a distribution include Hidden Markov Models (HMMs), and RNNs, or convolutional models.

### 3.4 Related Papers

I might use your help Sanmi here.

## 4 Experiments

### 4.1 Images

We learn generative models on the MNIST (hand written digits) and CELEB-A (celebrity faces). We compare our algorithm with VAE, standard GAN and Wasserstein GAN. As the main quality metric, we compare likelihoods computed on a test set using kernel density estimator (KDE).

For the MNIST dataset, we use an invertible perceptron in our approach to demonstrate that we can also use our approach to compute model likelihoods on the test set using the implicit generative model density function in Equation (6). (Note that in general our framework allows non-invertible mappings: We use a general convo-

lutional autoencoder for the CELEB-A dataset) The invertible perceptron we use for the MNIST dataset is defined as follows:

$$h_1 = \tanh_{\text{inv}}(\mathbf{Linear}[K, 600](h)), \\ x = \sigma_{\text{inv}}(\mathbf{Linear}[600, 784](h_1)),$$

where  $h \in \mathbb{R}^K$  denotes the latent representation, and  $\mathbf{Linear}[L_1, L_2](h) = Wh + b$ ,  $W \in \mathbb{R}^{L_2 \times L_1}$ ,  $b \in \mathbb{R}^{L_2}$  represents a linear layer (we follow the pytorch API convention to denote the input and output dimensionalities). The invertible non-linearity functions are denoted with  $\tanh_{\text{inv}}(\cdot)$ , and  $\sigma_{\text{inv}}(\cdot)$ , which respectively stand for invertible tangent-hyperbolic and invertible sigmoid functions. We basically use the original non linearity in the invertible regime, and a linear function in the saturation regimes. Namely, for hyperbolic tangent we have the following function:

$$\tanh_{\text{inv}}(t) = \begin{cases} ct - b & t \leq -1 + \epsilon \\ \tanh(t) & |t| \leq 1 - \epsilon \\ ct + b & t \geq 1 - \epsilon \end{cases}, \quad (9)$$

We use  $c = 0.01$ , and choose the bias term  $b$ , and the threshold  $\epsilon$  so that the function is continuous and smooth (has a continuous first derivative). Similarly, the invertible sigmoid function is defined as follows:

$$\sigma_{\text{inv}}(t) = \begin{cases} ct - b & t \leq \epsilon \\ \sigma(t) & 0 \leq t \leq 1 - \epsilon \\ ct + b & t \geq 1 - \epsilon \end{cases}, \quad (10)$$

Note that it is straightforward to derive the inverse functions once the parameters of the non-linearities are set. Therefore the inverse network is defined as follows:

$$h_1 = \sigma_{\text{inv}}^{-1}(\mathbf{Linear}^{-1}[784, 600](x)), \\ h = \tanh_{\text{inv}}^{-1}(\mathbf{Linear}^{-1}[600, K](h_1)),$$

where  $\mathbf{Linear}^{-1}[L_2, L_1](x) := (W^{\top}W)^{-1}W^{\top}(x - b)$ . Note that the parameters  $W, b$  are shared for a given forward and inverse Linear layers. To obtain the correction term due to the rectangular transformation, we consider Equation 5 which defined the implicit model pdf via CDF of the transformed density. We basically need to compute the volume change due to the transformation. (For square transformations this turns out to be the determinant of the Jacobian) We note that the volume change due to the rectangular linear transformation in a linear layer is given by  $\sqrt{\det(W^{\top}W)}$  [ ]. Therefore to the correction term involves dividing the original pdf with this volume change (we note that our transformation is approximately invertible due to the first step of our algorithm, and therefore we apply the inverse function theorem).

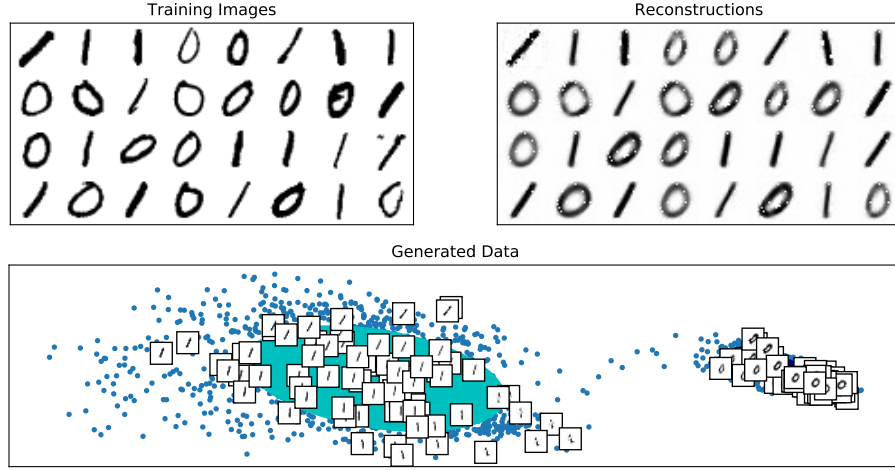


Figure 2: Demonstration of the two stage algorithm on a toy dataset with handwritten 0’s and 1’s. The purpose of this figure is to give a sense on how the proposed algorithm work. (Top row) Samples from the training set, and corresponding reconstructions. (Bottom row) Two dimensional embeddings of the training samples are shown with blue dots. We overlay sampled images. The solid color ellipses show the covariance components of the learned Gaussian mixture model for the base distribution.

To do objective comparisons between models we compute Kernel density estimates (KDE) on the test set: We sample 1000 points from the trained models, and represent the learned density as the sum of Kernel functions centered at these samples. We then compute the average score for all the test set. We use Gaussian Kernels, with bandwidth 0.01. Namely, the KDE scores we compute for the models are defined as follows:

$$\text{KDE score} = \frac{1}{N_{\text{test}} N_{\text{samples}}} \sum_{n=1}^{N_{\text{test}}} \sum_{m=1}^{N_{\text{samples}}} \mathcal{N}(x_n^{\text{test}}; x_m^{\text{sample}}, 0.1I).$$

Notice that for small kernel bandwidth, the above objective is tantamount to computing the nearest neighbor distance for all test instances. To get high scores from this estimator, the observed samples need to capture the diversity of the test instances. Also note that this estimator is computing an estimate for  $KL(p_{\text{test}} || p_{\text{model}})$ , so this metric penalizes mode collapse.

In the left panel of Figure 3, we compare the KDE scores for our two-stage algorithm, GAN, Wasserstein GAN and VAE on the MNIST dataset. We use the standard training-test split defined in the pytorch data utilities (60000 training instances and 6000 test instances). We try 7 different latent dimensionality  $K$  for all algorithms ranging from 20 to 140 with increments of 20. In our algorithm, we use a GMM with 30 full-covariance components for all  $K$  values. We see that performance drops

with increasing  $K$ , however we manage to stay better than VAEs and GANs. The performance drop is expected to happen with increasing  $K$ , because the density estimation problem in the latent space gets more difficult with increasing latent dimensionality. We would like to note that it is possible to use a more complicated base distribution and compensate.

In the right panel of Figure 3, we compare the model likelihood computed with the implicit likelihood equation in (6) with the base distribution likelihood (the complete likelihood minus the Jacobian term). The purpose of this is to examine if there is a correlation between these quantities. As we pointed out before, our algorithm does not require an exactly invertible mapping, and as can be seen from the figure the base distribution likelihood is somewhat correlated with the overall model likelihood, and therefore can potentially be used as a proxy for the complete likelihood for mappings for which we don’t know how to compute the Jacobian term.

In Figure ??, we show the random nearest neighbor samples for randomly selected test instances for all four algorithms in the top panel. We see that IML method is able to capture the diversity of the test instances well. On top of that we see much more definition in the generated images thanks to the multi-modal base distribution that we are using. As we earlier illustrated in Figure ??, using a simplistic base distribution causes a mismatch between

the mappings to the latent space and the draws from the base distribution. Due to the simplistic distributions used in VAEs, and GANs we see that these approaches tend to generate more samples which do not resemble handwritten digits. We also observe that quality of the samples (and nearest neighbor samples) are correlated with the KDE metric.

In Figure 4, we do the same nearest neighbor sample measurement on the CELEB-A dataset. We have set the latent dimensionality as 100 for all algorithms. We cropped the images using a face detector, and resized them to size  $64 \times 64$  in RGB space. We used 146209 such images for training, and 10000 images for test. We see that the proposed IML algorithm has more accurate nearest neighbor samples. We see that although the VAE is able to generate less distorted samples than DC-GAN, it's generated images contain more distortion than IML, potentially because of the simplistic latent representation.

For all algorithms we used the Adam optimizer [?]. As mentioned before, in the MNIST experiment, for IML we used the invertible network we introduced in this section. For GANs and VAE we used a standard one hidden layer perceptron with exact same sizes. Namely, the decoders maps  $K$  dimensions into 600, and 600 dimensions then gets mapped into 784 dimensions (MNIST images are of size  $28 \times 28$ ). We use the mirror image encoder for the VAE, that is we map 784 dimensions to 600, and that gets mapped into  $K$  dimensional vectors for the mean and variance of the posterior. For the CELEB-A dataset, we used a 5 layer convolutional encoders and decoders. (We used the basic DC-GAN generator architecture for all algorithms)




Figure 3: KDE likelihood with respect to the dimensionality of the latent space  $K$  on the MNIST dataset.

## 4.2 Audio

To show that our algorithm can be used to learn a generative model for sequential data, we experiment with generating speech and music. In all datasets, we work with audio with 8kHz sampling rate. We dissect the audio into

Table 1: KDE scores on test set for MNIST and CELEB-A datasets using 4 different algorithms

Algorithm	MNIST	CELEB-A
IML	143	-8183
VAE	132	-10941
GAN	-5	-11879
WGAN	51	

100ms long chunks, where consecutive chunks overlap by 50ms, and each window is multiplied by a Hann window. The autoencoder learns 80 dimensional latent representations for each chunk which is 800 samples long. We use three layer convolutional networks both in the encoder and decoder, where we use filters of 200 samples.

We fit an HMM to the extracted latent representations. We use 300 HMM states. The random samples are obtained by sampling from the fitted HMMs, and passing the sampled latent representation through the decoder. To reconstruct the generated chunks as an audio waveform, we follow the overlap-add procedure [ ]: We overlap the each generated by chunk by 50 percent and add.

As a speech experiment, we learn a generative model over digit utterances. We work with the free spoken digit dataset [?]. In Figure 6, we show the spectrograms of generated digit utterances along with spectrograms of the training digit utterances. Note that the generated digit utterances are generated in sequences (We generate one long sequence which contains multiple digits). We consider the cases where the training data only contains one digit type, and the case where the training data contains all digits. In the figure we show three case for the one-digit training task and the last figure shows the case in which we train the model jointly for all digits. We see that we are able to learn a generative model over one digit with a some variety.

As the music experiment, we train a model on a 2 minute long violin piece. We show the spectrogram of the first 10 seconds of the piece and our generated sequence in Figure ??.

## 5 Discussion

The algorithm we propose in this paper is very simple and effective. It is also principled in the sense that it performs maximum likelihood. We would like to emphasize that, compared the GANs the performance is much less sensitive to the network design choices and training parameters such as the learning rate. In author's experience, GANs are extremely sensitive to training param-






Figure 4: Samples from the MNIST dataset. **(top)** Generated nearest neighbor samples (nearest to test instances which are shown on the top row) for 4 different algorithms. **(bottom-left)** Random images generated with a VAE, **(bottom-right)** Generated Samples with our approach, samples from the same cluster are grouped together.



Figure 5: Samples from the CELEB-A dataset. **(top)** Generated nearest neighbor samples (nearest to test instances which are shown above) for 2 different algorithms. **(bottom-left)** Random images generated with GAN, **(bottom-right)** Generated Samples with our approach, samples from the same cluster are grouped together.



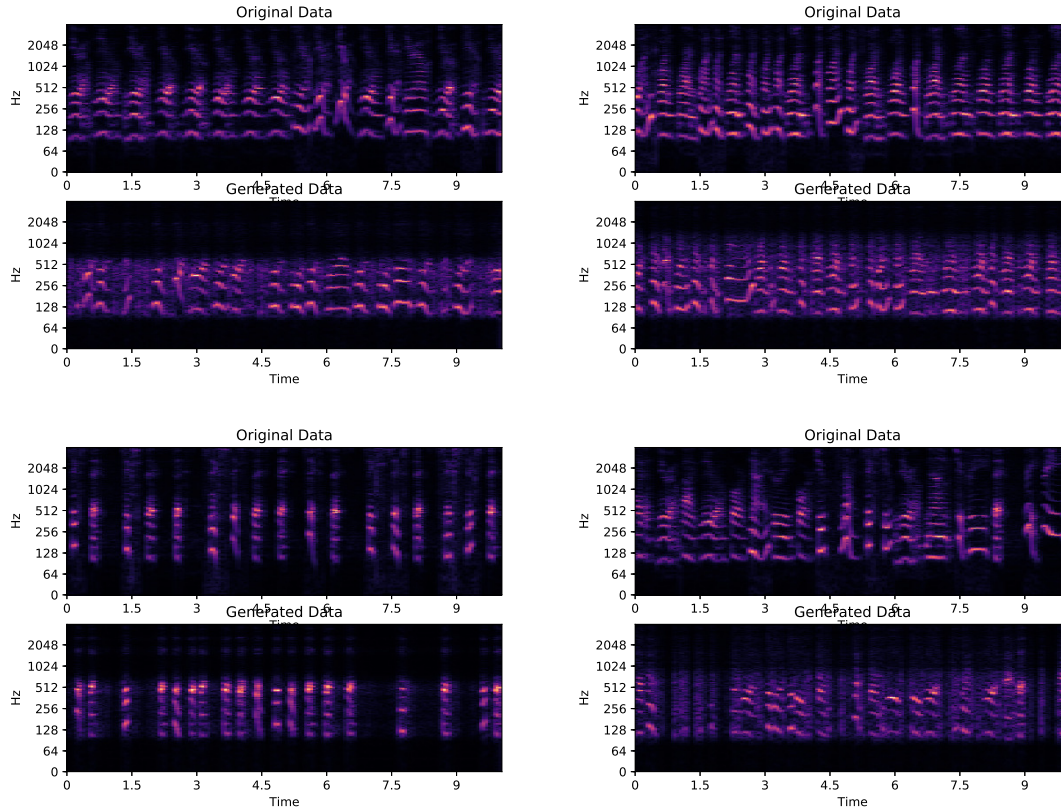
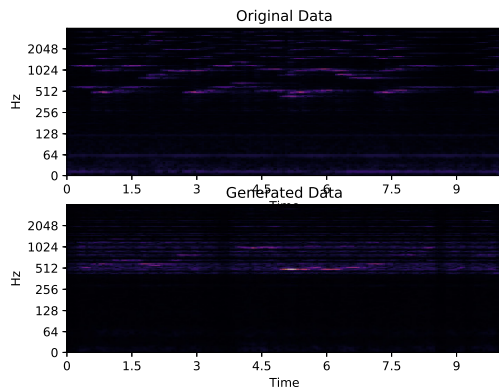


Figure 6: We illustrate the spectrograms for generated digits. **(top left)** Zero digits **(top right)** Generated sequence for digit 1 **(bottom left)** Generated sequence for digit 6, **bottom right** Generated sequence for all digits.



eters such as the learning rate. We have observed that decoupling the training of the base distribution from the neural network mapping makes the training much easier: In our approach it suffices to pick a small enough learning rate so that the encoder converges, and successfully embeds the data in a lower dimensional space.

In our experience, VAE's seem to be easier to train (much less susceptible to hyperparameter choices). However,

as we have seen in the results and figures, the simplistic choice for the base distribution results in distorted outputs. In our experiments we have used relatively more standard models to model the latent distribution, but it is possible to use complex methods such as Dirichlet Process Mixture models to obtain complicated base distributions.

## References

- [1] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *CoRR*, abs/1701.07875, 2017.
- [2] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.
- [3] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013.

- [4] S. Mohamed and B. Lakshminarayanan. Learning in Implicit Generative Models. *ArXiv e-prints*, October 2016.