

# MPCC: Matching Priors and Conditionals for Clustering

Nicolás Astorga<sup>1,4</sup>[0000–0001–6172–9362], Pablo Huijse<sup>2,4</sup>[0000–0003–3541–1697],  
Pavlos Protopapas<sup>3</sup>, and Pablo Estévez<sup>1,4</sup>[0000–0001–9164–4722]

<sup>1</sup> Department of Electrical Engineering, Universidad de Chile

<sup>2</sup> Informatics Institute, Faculty of Engineering Sciences, Universidad Austral de Chile

<sup>3</sup> Institute for Applied Computational Science, Harvard University

<sup>4</sup> Millennium Institute of Astrophysics, Chile

nicolas.astorga.r@ug.uchile.cl, phuijse@inf.uach.cl,

pavlos@seas.harvard.edu, pestevéz@ing.uchile.cl

**Abstract.** Clustering is a fundamental task in unsupervised learning that depends heavily on the data representation that is used. Deep generative models have appeared as a promising tool to learn informative low-dimensional data representations. We propose Matching Priors and Conditionals for Clustering (MPCC), a GAN-based model with an encoder to infer latent variables and cluster categories from data, and a flexible decoder to generate samples from a conditional latent space. With MPCC we demonstrate that a deep generative model can be competitive/superior against discriminative methods in clustering tasks surpassing the state of the art over a diverse set of benchmark datasets. Our experiments show that adding a learnable prior and augmenting the number of encoder updates improve the quality of the generated samples, obtaining an inception score of  $9.49 \pm 0.15$  and improving the Fréchet inception distance over the state of the art by a 46.9% in CIFAR10.

## 1 Introduction

Clustering is a fundamental unsupervised learning problem that aims to group the input data based on a similarity criterion. Traditionally, clustering models are trained on a transformed low-dimensional version of the original data obtained via feature-engineering or dimensionality reduction *e.g.* PCA. Hence the performance of clustering relies heavily on the quality of the feature space representation. In recent years deep generative models have been successful in learning low-dimensional representations from complex data distributions, and two particular models have gained wide attention: The Variational Autoencoder (VAE) [27], [47] and the Generative Adversarial Network (GAN) [14].

In VAE an encoder and a decoder network pair is trained to map the data to a low-dimensional latent space, and to reconstruct it back from the latent space, respectively. The encoder is used for inference while the decoder is used for generation. The main limitations of the standard VAE are the restrictive assumptions associated with the explicit distributions of the encoder and decoder

outputs. For the latter this translates empirically as loss of detail in the generator output. In GAN a generator network that samples from latent space is trained to mimic the underlying data distribution while a discriminator is trained to detect whether the generated samples are true or synthetic (fake). This adversarial training strategy avoids explicit assumptions on the distribution of the generator, allowing GANs to produce the most realistic synthetic outputs up to date [3], [25], [26]. The weaknesses of the standard GAN are the lack of inference capabilities and the difficulties associated with training (*e.g.* mode collapse).

One would like to combine the strengths of these two models, *i.e.* to be able to infer the latent variables directly from data and to have a flexible decoder that learns faithful data distributions. Additionally, we would like to train simultaneously for feature extraction and clustering as this performs better according to [59], [62], [61]. Extensions of the standard VAE that modify the prior distribution to make it suitable for clustering have been proposed in [24], [8], [23], although they still suffer from too restrictive generator models. On the other hand, the standard GAN has been extended to infer categories [52], [33], [6]. Other works have extended GAN to infer the posterior distribution of the latent variables reporting good results in both reconstruction and generation [5], [12], [9], [53], [34], [35]. These models do inference and have flexible generators but were not designed for clustering.

In this paper we propose a model able to learn good representations for clustering in latent space. The model is called Matching Priors and Conditionals for Clustering (MPCC). This is a GAN-based model with (a) a learnable mixture of distributions as prior for the generator, (b) an encoder to infer the latent variables from the data and (c) a clustering network to infer the cluster membership from the latent variables. Code are available at [github.com/jumpynitro/MPCC](https://github.com/jumpynitro/MPCC).

## 2 Background

MPCC is based on a matching joint distribution optimization framework. Let us denote  $q(x)$  as the true distribution and  $p(z)$  the prior, where  $x \in \mathcal{X}$  is the observed variable and  $z \in \mathcal{Z}$  is the latent variable, respectively.  $q(x)$  and  $p(z)$  stand for the marginalization of the inference model  $q(x, z)$  and generative model  $p(x, z)$ , respectively. If the joint distributions  $q(x, z)$  and  $p(x, z)$  match then it is guaranteed that all the conditionals and marginals also match. Intuitively this means that we can reach one domain starting from the other, *i.e.*, we have an encoder that allows us to reach the latent variables  $p(z) \approx q(z) = \mathbb{E}_{q(x)}[q(z|x)]$  and a generator that approximates the real distribution  $q(x) \approx p(x) = \mathbb{E}_{p(z)}[p(x|z)]$ . Notice that the latter approximation corresponds to a GAN optimization problem. In the case of vanilla GAN the Jensen-Shannon divergence  $D_{JS}(q(x)||p(x))$  is minimized, but other distances can be used [1], [16], [45], [41].

Although other classifications can be done [65], we recognize that the joint distribution matching problem can be divided in three general categories: i) matching the joints directly, ii) matching conditionals in  $\mathcal{Z}$  and marginals in  $\mathcal{X}$ , and iii) matching conditionals in  $\mathcal{X}$  and marginals in  $\mathcal{Z}$ . The straight forward

approach is to minimize the distance between the joint distributions using a fully adversarial optimization such as [12], [9], [10], which yields competitive results but still shows difficulties in reconstruction tasks likely affecting unsupervised representation learning. According to [34] these issues are related to the lack of an explicit optimization of the conditional distributions.

Recent works [49], [35], [65] have shown that the VAE [27] loss function (ELBO) is related to matching the inference and generative joint distributions. This can be demonstrated for the Kullback–Leibler (KL) divergence of  $p$  from  $q$ , which we refer as forward KL, as follows:

$$\begin{aligned}
 & D_{KL}(q(z, x)||p(z, x)) \\
 &= \mathbb{E}_{q(x)}[D_{KL}(q(z|x)||p(z|x))] + D_{KL}(q(x)||p(x)) \\
 &= \mathbb{E}_{q(x)}\mathbb{E}_{q(z|x)}[-\log p(x|z)] + \mathbb{E}_{q(x)}[D_{KL}(q(z|x)||p(z))] + \mathbb{E}_{q(x)}[\log q(x)] \\
 &= \mathbb{E}_{q(x)}[-\text{ELBO}] + \mathbb{E}_{q(x)}[\log q(x)], \tag{1}
 \end{aligned}$$

hence maximizing the ELBO can be seen as matching the conditionals in latent space  $\mathcal{Z}$  and the marginals in data space  $\mathcal{X}$  (see the second line in Eq. 1). The proof for the first equivalence in Eq. 1 can be found in the Appendix A.

In order to avoid latent collapse and the parametric assumptions of VAE, AIM [35] proposed the opposite, *i.e.* to match the conditionals in data space and the marginals in latent space. Starting from the KL divergence of  $q$  from  $p$ , which we refer as reverse KL, they obtained the following:

$$\begin{aligned}
 & D_{KL}(p(z, x)||q(z, x)) \\
 &= \mathbb{E}_{p(z)}[D_{KL}(p(x|z)||q(x|z))] + D_{KL}(p(z)||q(z)) \\
 &= \mathbb{E}_{p(z)}\mathbb{E}_{p(x|z)}[-\log q(z|x)] + \mathbb{E}_{p(z)}[D_{KL}(p(x|z)||q(x))] + \mathbb{E}_{p(z)}[\log p(z)], \tag{2}
 \end{aligned}$$

where  $p(z)$  is a fixed parametric distribution hence  $\mathbb{E}_{p(z)}[\log p(z)]$  is constant. Therefore [35] achieves the matching of joint distributions by minimizing  $D_{KL}(p(x|z)||q(x))$  to learn the real domain, and maximizing the likelihood of the encoder  $\mathbb{E}_{p(x|z)}[\log q(z|x)]$ . This allows obtaining an overall better performance than [12], [9], [34] in terms of reconstruction and generation scores. This method matches the conditional distribution explicitly, uses a flexible generator [14] and avoids latent collapse problems [38].

Lot of research has been done in unsupervised and semi supervised learning using straight forward joint distribution optimization [9], [12], [10], [21], and even more for conditional in latent space decomposition [27], [28], [40], [39]. In this work we explore the representation capabilities of the decomposition proposed in [35]. Our main contributions are:

- A mathematical derivation that allows us to have a varied mixture of distributions in latent space enforcing its clustering capabilities. Based on this derivation we developed a new generative model for clustering called MPCC, trained by matching prior and conditional distributions jointly.
- A comparison with the state-of-the-art showing that MPCC outperforms generative and discriminative models in terms of clustering accuracy and generation quality.

- An ablation study of the most relevant parameters of MPCC and a comparison with the AIM baseline [35] using state of the art architectures [10].

### 3 Method

#### 3.1 Model definition

MPCC extends the usual joint distribution of variables  $x \in \mathcal{X}$  and  $z \in \mathcal{Z}$  incorporating an additional latent variable,  $y \in \mathcal{Y}$ , which represents a given cluster. We specify the graphical models for *generation* and *inference* as

- $p(x, z, y) = p(y)p(z|y)p(x|z, y)$ ,
- $q(x, z, y) = q(y|z)q(z|x)q(x)$ ,

respectively. The only assumption in the graphical model is  $q(y|z) = q(y|z, x)$ , *i.e.*  $z$  contains all the information from  $x$  that is necessary to estimate  $y$ .

For generation, we seek to match the decoder  $p(x|z, y)$  to the real data distribution  $q(x)$ . The latent variable is defined by the conditional distributions  $p(z|y)$  which in general can be any distribution under certain restrictions (Section 3.3). The marginal distribution  $p(y)$  is defined as multinomial with weight probabilities  $\phi$ . Note that under this graphical model the latent space becomes multimodal defined by a mixture of distributions  $p(z) = \sum_y p(y)p(z|y)$ .

In the inference procedure the latent variables are obtained by the conditional posterior  $q(z|x)$  using the empirical data distribution  $q(x)$ . The distribution  $q(y|z)$  is a posterior approximation of the cluster membership of the data.

We call our model Matching Priors and Conditionals for Clustering (MPCC) and we optimize it by minimizing the reverse Kullback-Leibler divergence of the conditionals and priors between the inference and generative networks as follows:

$$\begin{aligned} & D_{KL}(p(x, z, y)||q(x, z, y)) \\ &= \mathbb{E}_{p(z, y)}[D_{KL}(p(x|z, y)||q(x|z, y))] \\ &+ \mathbb{E}_{p(y)}[D_{KL}(p(z|y)||q(z|y))] + D_{KL}(p(y)||q(y)). \end{aligned} \quad (3)$$

The proof for Eq. (3) can be found in Appendix A. In the following sections we derive a tractable expression for Eq. (3) and present the MPCC algorithm.

#### 3.2 Loss function

Because  $q(y)$ ,  $q(z|y)$  and  $q(x|z, y)$  are impossible to sample from, we derive a closed-form solution for Eq. (3). In particular for any fixed  $y$  and  $z$  we can decompose  $D_{KL}(p(x|z, y)||q(x|z, y))$  as follows:

$$\begin{aligned} & D_{KL}(p(x|z, y)||q(x|z, y)) \\ &= \mathbb{E}_{p(x|z, y)} \left[ \log \frac{p(x|z, y)}{q(x)} \frac{q(z, y)}{q(z, y|x)} \right] \\ &= \mathbb{E}_{p(x|z, y)} \left[ \log \frac{p(x|z, y)}{q(x)} - \log q(y|z) - \log q(z|x) + \log q(z|y) + \log q(y) \right]. \end{aligned} \quad (4)$$

Adding  $\log p(z|y) + \log p(y) - \log q(z|y) - \log q(y)$  to both sides of Eq. (4) and taking the expectation with respect to  $p(z, y)$  the Eq. (3) is recovered. After adding these terms and taking the expectation we can collect the resulting right hand side of Eq. (4) as follows:

$$\begin{aligned}
& \mathbb{E}_{p(z,y)}[D_{KL}(p(x|z,y)||q(x|z,y)) + D_{KL}(p(z|y)||q(z|y)) + D_{KL}(p(y)||q(y))] \\
= & \underbrace{\mathbb{E}_{p(y)p(z|y)}[D_{KL}(p(x|z,y)||q(x))]}_{\text{Loss I}} + \underbrace{\mathbb{E}_{p(y)p(z|y)p(x|z,y)}[-\log q(z|x) - \log q(y|z)]}_{\text{Loss II}} \\
& + \underbrace{\mathbb{E}_{p(z|y)p(y)}[\log p(y) + \log p(z|y)]}_{\text{Loss III}}, \tag{5}
\end{aligned}$$

where **Loss I** seeks to match the true distribution  $q(x)$ , **Loss II** is related to the variational approximation of the latent variables and **Loss III** is associated with the distribution of the cluster parameters. The right hand term of Eq. (5) is a loss function, composed of three terms with distributions that we can sample from. In the next section we explain the strategy to optimize each of the terms of the proposed loss function.

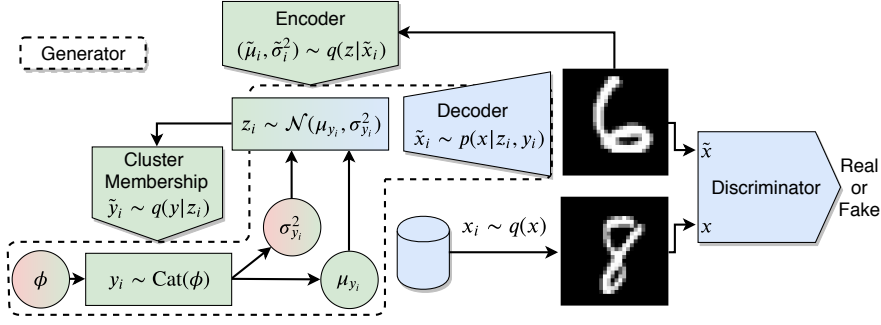
MPCC follows the idea that the data space  $\mathcal{X}$  is compressed in the latent space  $\mathcal{Z}$  and a separation in this space will likely partition the data in the most representatives clusters  $p(z|y)$ . The separability of these conditional distributions will be enforced by  $q(y|z)$  which also backpropagates through the parameters of  $p(z|y)$ . The connection with the data space is through the decoder  $p(x|z, y)$  for generation and the encoder  $q(z|x)$  for inference.

### 3.3 Optimizing MPCC

In what follows we describe the assumptions made in the distributions of the graphical model and how to optimize Eq. 5. For simplicity we assume the conditional  $p(z|y)$  to be a Gaussian distribution, but other distributions could be used with the only restriction being that their entropy should have a closed-form or at least a bound (second term in **Loss III**). In our experiments the latent variable  $z|y \sim \mathcal{N}(\mu_y, \sigma_y^2)$  is sampled using the reparameterization trick [27], *i.e.*  $z = \mu_y + \sigma_y \odot \epsilon$  where  $\epsilon \sim \mathcal{N}(0, I)$  and  $\odot$  is the Hadamard product. The parameters  $\mu_y, \sigma_y^2$  are learnable and they are conditioned on  $y$ . Under Gaussian conditional distribution the latent space becomes a GMM, as we can observe mathematically  $p(z) = \sum_y p(y)p(z|y) = \sum_y p(y)\mathcal{N}(\mu_y, \sigma_y^2)$ .

The distribution  $p(x|z, y)$  is modeled by a neural network and trained via adversarial learning, *i.e.* it does not require parametric assumptions. The inferential distribution  $q(z|x)$  is also modeled by a neural network and its distribution is assumed Gaussian for simplicity. The categorical distribution  $q(y|z)$  may also be modeled by a neural network but we propose a simpler approach based on the membership from the latent variable  $z$  to the Gaussian components. A diagram of the proposed model considering these assumptions is shown in Fig. 1. We now expand on this for each of the losses in Eq. 5.

**Loss I:** Instead of minimizing the Kullback-Leibler divergence shown in the first term on the right hand of Eq. (5) we choose to match the conditional decoder



**Fig. 1.** Diagram of the MPCC model. The blue colored elements are associated with **Loss I** (Eq. 6). The green colored elements are associated with **Loss II** (Equations 7 and 8). The red colored elements are associated with **Loss III** (Eq. 10). The dashed line corresponds to the generator (GMM plus decoder).

$p(x|z, y)$  with the empirical data distribution  $q(x)$  using a generative adversarial approach. The GAN loss function can be formulated as [11]

$$\begin{aligned} \max_D \mathbb{E}_{x \sim q(x)} [f(D(x))] + \mathbb{E}_{\tilde{x} \sim p(x, z, y)} [g(D(\tilde{x}))], \\ \min_G \mathbb{E}_{\tilde{x} \sim p(x, z, y)} [h(D(\tilde{x}))], \end{aligned} \quad (6)$$

where  $D$  and  $G$  are the discriminator and generator networks, respectively, and tilde is used to denote sampled variables. For all our experiments we use the hinge loss function [36], [55], *i.e.*  $f = -\min(0, o - 1)$ ,  $g = \min(0, -o - 1)$  and  $h = -o$ , being  $o$  the output of the discriminator. The parameters and distribution associated with **Loss I** are colored in blue in Fig.1.

**Loss II:** The first term of this loss is estimated through Monte Carlo sampling as

$$\begin{aligned} & \mathbb{E}_{p(y)p(z|y)p(x|z,y)} [-\log q(z|x)] \\ &= \mathbb{E}_{y_i \sim p(y), z_i \sim p(z|y=y_i), \tilde{x}_i \sim p(x|z=z_i, y=y_i)} \underbrace{\left[ \sum_{j=1}^J \frac{1}{2} \log(2\pi\tilde{\sigma}_{ij}^2) + \frac{(z_{ij} - \tilde{\mu}_{ij})^2}{2\tilde{\sigma}_{ij}^2} \right]}_{L_q(\tilde{\mu}_i, \tilde{\sigma}_i^2, z_i)}, \end{aligned} \quad (7)$$

where  $J$  is the dimensionality of the latent variable  $z$ . By minimizing Eq. (7) we are maximizing the log-likelihood of the encoder  $q(z|x)$  with respect to the Gaussian prior  $p(z|y)$ . This reconstruction error is estimated by matching the samples  $z_i \sim p(z|y = y_i)$  with the Gaussian distribution  $(\tilde{\mu}_i, \tilde{\sigma}_i^2) \sim q(z|x = \tilde{x}_i)$ , where  $\tilde{x}_i$  is the decoded representation of  $z_i$ .

The second term of **Loss II** is equivalent to the cross-entropy between the sampled label  $y_i \sim p(y)$  and the estimated cluster membership  $\tilde{y}_i$

$$L_c(y_i, \tilde{y}_i) = - \sum_{k=1}^K y_{ik} \log \tilde{y}_{ik}, \quad (8)$$

where  $K$  is the number of clusters and

$$\tilde{y}_{im} = q(y = m | z = z_i) = \frac{\mathcal{N}(z_i | \mu_m, \sigma_m^2)}{\sum_{k=1}^K \mathcal{N}(z_i | \mu_k, \sigma_k^2)}, \quad (9)$$

is the membership of  $z_i$  to the  $m$ -th cluster. The parameters  $\mu_m$  and  $\sigma_m^2$  are learnable, and  $m \in [1, \dots, K]$  is the index corresponding to each cluster. The parameters and distribution associated with **Loss II** are colored in green in Fig. 1. In practice Eq. 9 is estimated using the log-sum-exp trick.

**Loss III:** This loss is associated with the regularization of the Gaussian mixture model parameters  $\phi$ ,  $\mu$  and  $\sigma^2$  and has a closed form

$$\begin{aligned} & \mathbb{E}_{p(y)p(z|y)}[\log p(y) + \log p(z|y)] \\ &= \underbrace{\sum_{k=1}^K \phi_k \left[ \log \phi_k - \sum_{j=1}^J \left( \frac{1}{2} + \frac{1}{2} \log(2\pi\sigma_{kj}^2) \right) \right]}_{L_p(\phi, \sigma^2)}, \end{aligned} \quad (10)$$

where the first term corresponds to the entropy maximization of the mixture weights, *i.e.* in general every Gaussian will not collapse to less than  $K$  modes of the data distribution which is a solution with lower entropy. In our experiments we fix  $\phi_k = 1/K$ , *i.e.*  $\phi$  is not learnable. The second term is a regularization for the variance (entropy) of each Gaussian which avoids the collapse of  $p(z|y)$ . The parameters associated with **Loss III** are shown in red in Fig. 1.

**Loss I** scale differs from that of the terms associated with the latent variables. To balance all terms we multiply Eq. (7) by one over the dimensionality of  $x^5$  and the second term of Eq. (10) by one over the dimensionality of the latent variables. During training **Loss III** is weighted by a constant factor  $\lambda_p$ . We explain how this constant is set in Section 5.3. The full procedure to train the MPCC model is summarized in Algorithm 1. Note that MPCC is scalable in the number of clusters since Eq. 7 is a Monte Carlo approximation in  $y$  and the cost of Eq. 10 is low since  $J$  is small in comparison to the data dimensionality.

## 4 Related methods

In Section 3 we showed that the latent space of MPCC it is reduced to a GMM under Gaussian conditional distribution. Because all the experiments are performed based on this assumption in this section we summarize the literature of generative and autoencoding models that consider GMMs.

<sup>5</sup> If  $x$  is an image then its dimensionality would be *channels*  $\times$  *height*  $\times$  *width*

**Algorithm 1** MPCC algorithm

---

```

1:  $K, J \leftarrow$  Set number of clusters and latent dimensionality
2:  $\eta, \eta_p \leftarrow$  Set learning rates
3:  $\theta_g, \theta_d, \theta_e \leftarrow$  Initialize network parameters
4:  $\phi, \mu, \sigma^2 \leftarrow$  Initialize GMM parameters
5:  $\theta_c \leftarrow [\phi, \mu, \sigma^2]$ 
6: repeat
7:   for  $D_{steps}$  do
8:      $x_1, \dots, x_n \sim q(x)$  ▷ Draw n samples from empirical distribution
9:      $y_1, \dots, y_n \sim p(y)$  ▷ Draw n samples from categorical prior
10:     $z_i \sim p(z|y = y_i), \quad i = 1, \dots, n$  ▷ Draw n samples from Gaussian conditional prior
11:     $\tilde{x}_i \sim p(x|z = z_i, y = y_i), \quad i = 1, \dots, n$  ▷ Generate samples using generator network
12:     $\theta_d \leftarrow \theta_d + \eta \nabla_{\theta_d} \left[ \frac{1}{n} \sum_{j=1}^n f(D(x_j)) + \frac{1}{n} \sum_{i=1}^n g(D(\tilde{x}_i)) \right]$  ▷ Gradient update on discriminator network
13:  end for
14:   $y_1, \dots, y_n \sim p(y)$  ▷ Draw n samples from categorical prior
15:   $z_i \sim p(z|y = y_i), \quad i = 1, \dots, n$  ▷ Draw n samples from Gaussian conditional prior
16:   $\tilde{x}_i \sim p(x|z = z_i, y = y_i), \quad i = 1, \dots, n$  ▷ Generate samples using generator network
17:   $(\theta_g, \theta_c) \leftarrow (\theta_g, \theta_c) - \eta \nabla_{(\theta_g, \theta_c)} \frac{1}{n} \sum_{i=1}^n h(D(\tilde{x}_i))$  ▷ Gradient update on generator network
18:  for  $E_{steps}$  do
19:     $y_1, \dots, y_n \sim p(y)$  ▷ Draw n samples from categorical prior
20:     $z_i \sim p(z|y = y_i), \quad i = 1, \dots, n$  ▷ Draw n samples from Gaussian conditional prior
21:     $\tilde{x}_i \sim p(x|z = z_i, y = y_i), \quad i = 1, \dots, n$  ▷ Generate samples using generator network
22:     $(\hat{\mu}_i, \hat{\sigma}_i^2) \sim q(z|x = \tilde{x}_i), \quad i = 1, \dots, n$  ▷ Encode  $\tilde{x}$  to obtain mean and variance
23:     $\theta_e \leftarrow \theta_e - \eta \nabla_{\theta_e} \frac{1}{n} \sum_{i=1}^n L_q(\hat{\mu}_i, \hat{\sigma}_i^2, z_i)$  ▷ Gradient update on encoder network
24:    if first  $E_{step}$  then
25:       $\tilde{y}_i \sim q(y|z = z_i), \quad i = 1, \dots, n$ 
26:       $\theta_c \leftarrow \theta_c - \eta_p \nabla_{\theta_c} \left[ \frac{1}{n} \sum_{i=1}^n L_c(y_i, \tilde{y}_i) + \lambda_p \cdot L_p(\phi, \sigma^2) \right]$  ▷ Gradient update on Prior parameters
27:    end if
28:  end for
29: until convergence

```

---

The combination of generative models and GMMs is not new. Several methods have applied GMM in autoencoding [56], [66] or GAN [17], [48] applications without clustering purposes. Other approaches have performed clustering but are not directly comparable since they use mixtures of various generators and discriminators [64] or fixed priors with ad-hoc set parameters [2].

Among the related works on generative models for clustering the closest approaches to our proposal are ClusterGAN [43] and Variational Deep Embedding [23]. ClusterGAN differs from our model in that it sets the dimensions of the latent space as either continuous or categorical while MPCC uses a continuous latent space which is conditioned on the categorical variable  $y$ . On the other hand, Variational Deep Embedding (VADE) differs greatly in the training procedure, despite its similar theoretical basis. VADE, as a variational autoencoder model, matches the joint distributions in the forward KL sense  $D_{KL}(q(x, z, y)||p(x, z, y))$  by matching the posteriors and the marginals in data space as demonstrated in Appendix B. MPCC optimizes the reverse KL, *i.e.* matching the priors in latent space and conditionals in data space. Optimizing different KLs yield notably different decompositions and thus training procedures. For the forward KL [23] in addition to the challenges in scaling to larger dimension (Section 2) it is more difficult to generalize the latent space to any multi-modal distribution, we briefly discuss the reasons for this in Appendix B.



## 5 Experiments

### 5.1 Quantitative Comparison

Following [59], the performance of MPCC is measured using the clustering accuracy metric in which each cluster is assigned to the most frequent class in the cluster. Formally this is defined as

$$\text{ACC} = \max_{m \in \mathcal{M}} \frac{\sum_{i=1}^N \mathbb{1}\{y_i = m(c_i)\}}{N}, \quad (11)$$

where  $N$  is the total number of samples,  $y_i$  is the ground truth,  $c_i = \arg \max_k q(y = k|z = z_i)$  is the predicted cluster and  $\mathcal{M}$  is the space of all possible mappings between clusters and labels.

To measure the quality of the samples generated by MPCC we use the inception score (IS) [50] and the Fréchet inception distance (FID) [19].

### 5.2 Datasets

In order to evaluate MPCC we performed clustering in five benchmark datasets: a handwritten digit dataset (MNIST, [32]), a handwritten character dataset (Omniglot, [31]), two color image dataset (CIFAR-10 and CIFAR-100 [30]) and a fashion products image dataset (Fashion-MNIST, [58]). For CIFAR-100 we consider the 20 superclasses. Omniglot was created using the procedure described in [20]. Because the task is fully unsupervised we concatenate the training and test sets as frequently done in the area [20], [4], [59]. All datasets have 10 classes except for Omniglot and CIFAR-20 with 100 and 20 respectively. All images were resized to  $32 \times 32$  and rescaled to  $[-1,1]$  in order to use similar architectures. The CIFAR-10 experiments where IS and FID are reported (tables 1, 2 and 3) were trained using only the training set (50,000 examples) for a fair comparison with the literature. For all clustering experiments we use the same number of cluster as the datasets classes.

### 5.3 Empirical details

Our architecture is based on optimization techniques used in the BigGAN [3]<sup>6</sup>, we found that simpler architectures such as DCGAN [46] were not able to learn complex distributions like CIFAR-10 while optimizing the parameters of the prior. Architecture details are given in Appendix C. We consider parameter sharing between the encoder and discriminator, and we test the importance of this in Section 5.5. We set  $D_{steps} = 4$  (see Algorithm 1), we found that using smaller values of  $D_{steps}$  causes mode collapse problems when training on CIFAR-10 (see Appendix D). A similar effect can be observed when choosing a low number of latent dimensions, therefore we set  $J = 128$  in all CIFAR-10 experiments. We made small changes in the architecture and optimization parameters depending on the dataset (see Appendix C).

<sup>6</sup> <https://github.com/ajbrock/BigGAN-PyTorch>

We observed the same relation between batch size and (IS, FID) reported in [3]. However we found artifacts that hurt accuracy performance when using batch size larger than 50. For simplicity we used this value in all experiments.

We consider a weighting factor  $\lambda_p$  for **Loss III** (Eq. 10). We observed that if  $\lambda_p = 1$ , the standard deviation of the prior  $\sigma$  would increase monotonically, hindering training. On the other hand if  $\lambda_p$  is too small,  $\sigma$  decreases, collapsing at some point. We found empirically that a value of  $\lambda_p = 0.01$  combined with a minimum threshold for  $\sigma$  of 0.5 allow the algorithm to converge to good solutions.

The parameter settings indicated above were fixed for all experiments and didn't show a big effect in accuracy performance. In section 5.4 we explore the parameters that most affect the training. We trained all experiments for 75,000 iterations, except for MNIST and Omniglot which iterate for 125,000. For unconditional and conditional training we kept the model of the last iteration.

#### 5.4 Ablation study

We found that  $E_{steps}$ , the number of encoder updates per epoch, and  $\eta_p$ , the learning rate of the prior parameters, are the most relevant hyperparameters to obtain high accuracy and generation quality. Increasing  $E_{steps}$  improves the estimation of  $q(z|x)$  since the prior and generator parameters are changing constantly. Rows 1-3 of Table 1 show that the reconstruction error (MSE) decreases with  $E_{steps}$ . Generation quality metrics (IS, FID) also improve with larger values of  $E_{steps}$  due to the shared parameters between encoder and discriminator.

At initialization the GMM components might not be separated. We observed that the clustering accuracy drops when the generators learns a good approximation of the real distribution before the clusters are separated. To avoid this we use a larger learning rate for the parameters of the GMM prior with respect to the parameters of the generator, encoder and discriminator. Rows 4-6 of Table 1 show that the clustering accuracy increases for larger values of  $\eta_p$ .

#### 5.5 Comparison between GMM Prior and Normal Prior

Using the best configuration found in the ablation study, we performed a comparison with AIM [35], whose results are shown in Table 2. We can consider AIM as a particular case of MPCC where a standard Normal prior is used instead of the GMM prior. AIM does not perform clustering therefore we compare it with MPCC in terms of reconstruction and generation quality. We use the same architecture and parameter settings of MPCC for AIM and we denote this model as AIM-MPCC. To extend our analysis further, Table 2 includes the results of using parameter sharing between the encoder and the discriminator (Appendix C), an idea that was considered but not fully explored in [35].

Note that AIM-MPCC (NS) is considered a baseline because the prior is Gaussian and the encoder doesn't share parameters with the discriminator thus the existence of the encoder doesn't affect the generation quality. In Table 2 we can observe the relevance of parameter sharing, with this configuration ( $E_{steps} = 4$ ) the baseline improves by 0.85 (IS) and 11.13 (FID) points. Adding the GMM

**Table 1.**  $E_{steps}$  correspond to the encoder updates and  $\eta_p$  to the learning rate of the prior parameters. The scale of MSE is in  $10^{-3}$ . The statistics were obtained for at least three runs

$E_{steps}$	$\eta_p$	Acc %	IS	FID	MSE
1	2e-4	41.31	8.82	11.38	1.34
		$\pm 5.74$	$\pm 0.07$	$\pm 0.23$	$\pm 0.96$
2	2e-4	38.67	9.02	9.66	1.01
		$\pm 3.52$	$\pm 0.05$	$\pm 3.98$	$\pm 1.11$
4	2e-4	38.27	9.25	7.50	0.331
		$\pm 2.46$	$\pm 0.09$	$\pm 0.43$	$\pm 0.09$
4	4e-4	52.58	9.44	6.55	0.48
		$\pm 5.30$	$\pm 0.06$	$\pm 0.33$	$\pm 0.22$
4	6e-4	61.99	9.49	6.59	1.04
		$\pm 4.96$	$\pm 0.15$	$\pm 0.45$	$\pm 1.03$

**Table 2.** Comparison of MPCC and AIM-MPCC methods with sharing parameters (S) and without sharing (NS) on the CIFAR-10 dataset. The scale of MSE is in  $10^{-3}$ . The statistics were obtained for five runs

Model	Acc %	IS	FID	MSE
AIM-MPCC (NS)	-	8.24	21.55	1.52
		$\pm 0.07$	$\pm 1.47$	$\pm 0.84$
AIM-MPCC (S)	-	9.09	10.42	1.64
		$\pm 0.04$	$\pm 0.36$	$\pm 1.42$
MPCC (S)	61.99	9.49	6.59	1.04
	$\pm 4.96$	$\pm 0.15$	$\pm 0.45$	$\pm 1.03$

in the prior improves an additional 0.4 (IS) and 3.93 (FID) points. In total when using the GMM Prior and parameter sharing with additional encoder updates we improve the baseline from 21.55 to 6.59 (69.4% improvement) in terms of FID score and 1.25 points (15.2% improvement) in terms of IS. It is important to notice that these techniques are general and easily applied to any GAN scheme.

## 5.6 Generation quality of MPCC

Using the configuration of row five from Table 1 we compare MPCC with nine state of the art methods, surpassing them in terms of IS and FID scores in both the unsupervised and supervised setting, as shown in Table 3. The unconditional generation is the most significant with an improvement of 46.9% (FID) over state-of-the-art (SOTA), AutoGAN [13]. Most notably its performance is better than the current best conditional method (BigGAN).

## 5.7 Clustering experiments

Table 4 shows the clustering results for the selected benchmarks. We observe that in all the available benchmarks MPCC outperform the related methods, VADE [23] and ClusterGAN [43]. In more complex datasets such as CIFAR10, MPCC notably surpass discriminative based models (*e.g.* [20], [22]) which are the most competitive methods in the current literature. For benchmarks with more classes the margin is even larger obtaining improvements over the SOTA of  $\sim 42\%$  and  $\sim 9.7\%$  points in Omniglot and CIFAR-20 respectively, demonstrating empirically the scalability of MPCC when using a high number of clusters.

It can be observed that for all datasets our proposed method either achieves or surpasses the SOTA in terms of clustering. Figures 2 and 3 show examples of generated and reconstructed images, respectively, using the MPCC model with the highest accuracy in the MNIST and CIFAR-10 datasets.

**Table 3.** Inception and FID scores for CIFAR-10, in unconditional and conditional training. Higher IS is better. Lower FID is better. †: Average of 10 runs. ‡: Best of many runs. ††: Average of 5 runs. Results without symbols are not specified

Model	IS	FID	Model	IS	FID
DCGAN [46]	6.64 ± 0.14	–	WGAN-GP [16]	8.42 ± 0.10	–
SN-GAN† [42]	8.22 ± 0.05	21.7 ± 2.1	SN-GAN † [42]	8.60 ± 0.08	17.5
AutoGAN [13]	8.55 ± 0.10	12.42	Splitting GAN ‡ [15]	8.87 ± 0.09	–
PG-GAN ‡ [25]	8.80 ± 0.05	–	CA-GAN † [44]	9.17 ± 0.13	–
NCSN [51]	8.91	25.32	BigGAN [3]	9.22	–
<b>MPCC††</b>	9.49 ± 0.15	6.59 ± 0.45	<b>MPCC††</b>	9.55 ± 0.08	5.69 ± 0.17

(a) Unconditional (unsupervised) generation

(b) Conditional (supervised) generation

## 6 Discussion

Our results show that MPCC achieves a superior performance with respect to the SOTA on both clustering and generation quality. We note that the current SOTA on unsupervised and semisupervised learning relies on consistency training [60] and/or data augmentation [22], *i.e.* techniques that are complementary to MPCC and could be used to further improve our results.

To the best of our knowledge MPCC is the first deep generative clustering model capable of dealing with more complex distributions such as CIFAR-10/20 and the first to report clustering accuracy on these datasets. Additionally we empirically prove the scalability of MPCC showing significant improvements in datasets with a larger number of classes, 20 in case of CIFAR-20 and 100 in case of Omniglot, such scalability has not been proven for the current literature on generative models [23], [43], [52], [63].

Our experiments show that MPCC’s key innovations: GMM prior, loss function and optimization scheme (*e.g.* extra encoder updates with parameter sharing) are not only relevant to achieve a good clustering accuracy but also allows us to obtain unprecedented results in terms of generation quality (Table 3). Which translates in improvements of 69.4% over the baseline (Table 2) and 46.9% over the SOTA (Table 3) in terms of FID score. We think that the exceptional generation capabilities of MPCC are related to the support that each cluster covers of the real domain. Since each cluster learns a subset of the real distribution the interpolation between two points within a cluster is smoother compared to the case where no latent separation exists. The latter is explained by the learnable shared features which exploit the similarities existing in a cluster and are not present in a fixed global prior (*i.e.* ALI, AIM).

The high generation quality can be appreciated in Fig. 2 (more samples in Appendix E), where many clusters sample consistently different classes. However we can still see some classes mixed in some clusters, for example in columns 7-8 with cats and dogs. MPCC also presents a competitive performance in terms of

**Table 4.** Clustering accuracies for several methods and datasets. All the results of CIFAR-20 dataset were extracted from [22], the results of IMSAT and DEC from [20], the results of InfoGAN and ClusterGAN from [43] and the remaining from their respective papers. †: average of 5 or more runs. ‡: best of 5 runs. §: best of 10 or more runs. ‖: best of 3 runs. Results without symbols are not specified

Methods	Datasets				
	MNIST	Onmigtot	FMNIST	CIFAR-10	CIFAR-20
DEC [59]	84.3 <sup>§</sup>	5.3 ± 0.3 <sup>†</sup>	–	46.9 ± 0.9 <sup>†</sup>	18.5
VADE [23]	94.46 <sup>§</sup>	–	–	–	–
InfoGAN [5]	89.0 <sup>‡</sup>	–	61.0 <sup>‡</sup>	–	–
ClusterGAN [43]	95.0 <sup>‡</sup>	–	63.0 <sup>‡</sup>	–	–
DAC [4]	97.75 <sup>‖</sup>	–	–	52.18 <sup>‖</sup>	23.8
IMSAT (VAT) [20]	98.4 ± 0.4 <sup>†</sup>	24.0 ± 0.9 <sup>†</sup>	–	45.6 ± 0.8 <sup>†</sup>	–
ADC [18]	98.7 ± 0.6 <sup>†</sup>	–	–	29.3 ± 1.5 <sup>†</sup>	16.0
SCAE [29]	98.5 ± 0.10 <sup>†</sup>	–	–	33.48 ± 0.3 <sup>†</sup>	–
IIC [22]	98.4 ± 0.65 <sup>†</sup>	–	–	57.6 ± 0.3 <sup>†</sup>	25.5 ± 0.46 <sup>†</sup>
MPCC (Five runs)	98.48 ± 0.52	65.87 ± 1.46	62.56 ± 4.16	64.25 ± 5.31	35.21 ± 1.69
MPCC (Best three runs)	98.76 ± 0.03	66.95 ± 0.62	64.99 ± 2.22	67.73 ± 2.50	36.51 ± 0.71

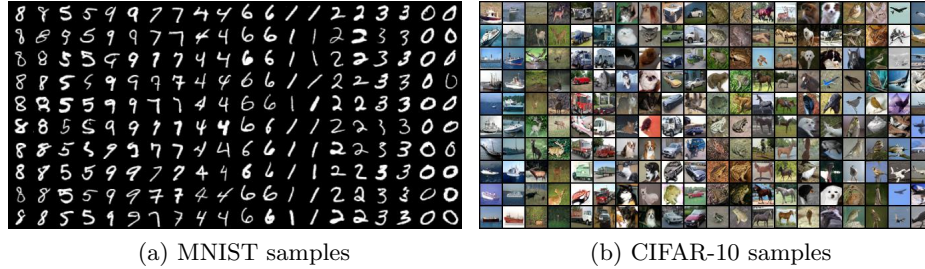
conditional distribution matching (Fig. 3). The errors observed in reconstruction are semantic and similar to those observed in [10].

MPCC opens the possibility of future research in many relevant topics which are out of the scope of this paper. Based on our experiments the most important extensions are: 1) Experiment with other conditional distributions  $p(z|y)$ , *e.g.* other exponential-family distributions or other flexible distributions by bounding their entropy (Section 3). This can be suitable for more expressive priors as it’s shown in recent work [57]. 2) Experiment with imbalanced distribution of classes by changing  $\phi$  accordingly, we consider this to be a relevant problem in the unsupervised setting which only a few works have addressed [54]. 3) Experiment with higher resolution datasets such as ImageNet [7] or CelebA [37]. Current works on clustering have not focus their attention to higher-resolution due to its complexity, MPCC is a promising approach to tackle this task from a semantic perspective [10].


## 7 Conclusions

We developed a new clustering algorithm called MPCC, derived from a joint distribution matching perspective with a latent space modeled by a mixture distribution. As a deep generative model this algorithm is suitable for interpretable representations, having both an inference and a generative network. The inference network allows us to infer the cluster membership and latent variables from data, while the generator performs sampling conditioned on the cluster category.

An important contribution of this work lays in the solid mathematical and optimization framework on which MPCC is based. This framework is general



**Fig. 2.** Generated images for a) MNIST and b) CIFAR-10 datasets, respectively. Every two columns we set a different value for the categorical latent variable  $y$ . *i.e.* the samples shown correspond to a different conditional latent space  $z \sim p(z|y)$ .



**Fig. 3.** Reconstructions for a) MNIST, and b) CIFAR-10 datasets, respectively. Odd columns represent real data and even columns correspond to their reconstructions.

and we recognize several opportunities to further enhance our model. The results obtained with MPCC improve over the state of the art in clustering methods. Most notably, our model is able to generate samples with an unprecedented high quality, surpassing the state of the art performances in both conditional and unconditional training in the CIFAR-10 dataset.

**Acknowledgement** Pablo Huijse and Pablo Estevez acknowledge financial support from the Chilean National Agency for Research and Development (ANID) through grants FONDECYT 1170305 and 1171678, respectively. Also, the authors acknowledge support from the ANID’s Millennium Science Initiative through grant IC12009, awarded to the Millennium Institute of Astrophysics, MAS. We acknowledge the InnovING 2030 project of the Faculty of Engineering Sciences UACH for the support and funding provided for the development of this work. We thank Pavlos Protopapas and the IACS for hosting Nicolás Astorga through the development of this research.

## References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: Precup, D., Teh, Y.W. (eds.) Proceedings of the 34th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 70, pp. 214–223. PMLR, International Convention Centre, Sydney, Australia (06–11 Aug 2017), <http://proceedings.mlr.press/v70/arjovsky17a.html>
2. Ben-Yosef, M., Weinshall, D.: Gaussian mixture generative adversarial networks for diverse datasets, and the unsupervised clustering of images. CoRR **abs/1808.10356** (2018), <http://arxiv.org/abs/1808.10356>
3. Brock, A., Donahue, J., Simonyan, K.: Large scale GAN training for high fidelity natural image synthesis. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=B1xsqj09Fm>
4. Chang, J., Wang, L., Meng, G., Xiang, S., Pan, C.: Deep adaptive image clustering. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 5880–5888 (Oct 2017). <https://doi.org/10.1109/ICCV.2017.626>
5. Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., Abbeel, P.: Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In: Advances in neural information processing systems. pp. 2172–2180 (2016)
6. Dai, Z., Yang, Z., Yang, F., Cohen, W.W., Salakhutdinov, R.R.: Good semi-supervised learning that requires a bad gan. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems 30, pp. 6510–6520. Curran Associates, Inc. (2017), <http://papers.nips.cc/paper/7229-good-semi-supervised-learning-that-requires-a-bad-gan.pdf>
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: CVPR09 (2009)
8. Dilokthanakul, N., Mediano, P.A., Garnelo, M., Lee, M.C., Salimbeni, H., Arulkumar, K., Shanahan, M.: Deep unsupervised clustering with gaussian mixture variational autoencoders. arXiv preprint arXiv:1611.02648 (2016)
9. Donahue, J., Krähenbühl, P., Darrell, T.: Adversarial feature learning. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings (2017), <https://openreview.net/forum?id=BJtNZAFgg>
10. Donahue, J., Simonyan, K.: Large scale adversarial representation learning. ArXiv **abs/1907.02544** (2019)
11. Dong, H.W., Yang, Y.H.: Towards a deeper understanding of adversarial losses. ArXiv **abs/1901.08753** (2019)
12. Dumoulin, V., Belghazi, I., Poole, B., Lamb, A., Arjovsky, M., Mastropietro, O., Courville, A.C.: Adversarially learned inference. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings (2017), <https://openreview.net/forum?id=B1ELR4cgg>
13. Gong, X., Chang, S., Jiang, Y., Wang, Z.: Autogan: Neural architecture search for generative adversarial networks. In: The IEEE International Conference on Computer Vision (ICCV) (Oct 2019)
14. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in neural information processing systems. pp. 2672–2680 (2014)

- 16 N. Astorga, P. Huijse, P. Protopapas, P. Estévez.
15. Grinblat, G.L., Uzal, L.C., Granitto, P.M.: Class-splitting generative adversarial networks (2017)
  16. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.: Improved training of wasserstein gans. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. pp. 5769–5779. NIPS’17, Curran Associates Inc., USA (2017), <http://dl.acm.org/citation.cfm?id=3295222.3295327>
  17. Gurumurthy, S., Kiran Sarvadevabhatla, R., Venkatesh Babu, R.: Deligan : Generative adversarial networks for diverse and limited data. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
  18. Häusser, P., Plapp, J., Golkov, V., Aljalbout, E., Cremers, D.: Associative deep clustering: Training a classification network with no labels. In: GCPR (2017)
  19. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems 30, pp. 6626–6637. Curran Associates, Inc. (2017)
  20. Hu, W., Miyato, T., Tokui, S., Matsumoto, E., Sugiyama, M.: Learning discrete representations via information maximizing self-augmented training. In: Proceedings of the 34th International Conference on Machine Learning - Volume 70. pp. 1558–1567. ICML’17, JMLR.org (2017), <http://dl.acm.org/citation.cfm?id=3305381.3305542>
  21. Jesson, A., Low-Kam, C., Nair, T., Soudan, F., Chandelier, F., Chapados, N.: Adversarially learned mixture model. CoRR [abs/1807.05344](https://arxiv.org/abs/1807.05344) (2018), <http://arxiv.org/abs/1807.05344>
  22. Ji, X., Henriques, J.F., Vedaldi, A.: Invariant information clustering for unsupervised image classification and segmentation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9865–9874 (2019)
  23. Jiang, Z., Zheng, Y., Tan, H., Tang, B., Zhou, H.: Variational deep embedding: an unsupervised and generative approach to clustering. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence. pp. 1965–1972. AAAI Press (2017)
  24. Johnson, M.J., Duvenaud, D.K., Wiltchko, A., Adams, R.P., Datta, S.R.: Composing graphical models with neural networks for structured representations and fast inference. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) Advances in Neural Information Processing Systems 29, pp. 2946–2954. Curran Associates, Inc. (2016)
  25. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=Hk99zCeAb>
  26. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. CoRR [abs/1812.04948](https://arxiv.org/abs/1812.04948) (2018), <http://arxiv.org/abs/1812.04948>
  27. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings (2014), <http://arxiv.org/abs/1312.6114>
  28. Kingma, D.P., Mohamed, S., Jimenez Rezende, D., Welling, M.: Semi-supervised learning with deep generative models. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.)



- Advances in Neural Information Processing Systems 27, pp. 3581–3589. Curran Associates, Inc. (2014), <http://papers.nips.cc/paper/5352-semi-supervised-learning-with-deep-generative-models.pdf>
29. Kosiorek, A., Sabour, S., Teh, Y.W., Hinton, G.E.: Stacked capsule autoencoders. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems 32, pp. 15512–15522. Curran Associates, Inc. (2019), <http://papers.nips.cc/paper/9684-stacked-capsule-autoencoders.pdf>
  30. Krizhevsky, A.: Learning multiple layers of features from tiny images. Tech. rep., Citeseer (2009)
  31. Lake, B.M., Salakhutdinov, R., Tenenbaum, J.B.: Human-level concept learning through probabilistic program induction. *Science* **350**(6266), 1332–1338 (2015)
  32. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
  33. Li, C., Xu, T., Zhu, J., Zhang, B.: Triple generative adversarial nets. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems 30, pp. 4088–4098. Curran Associates, Inc. (2017), <http://papers.nips.cc/paper/6997-triple-generative-adversarial-nets.pdf>
  34. Li, C., Liu, H., Chen, C., Pu, Y., Chen, L., Heno, R., Carin, L.: Alice: Towards understanding adversarial learning for joint distribution matching. In: Advances in Neural Information Processing Systems. pp. 5495–5503 (2017)
  35. Li, H., Wang, Y., Chen, C., Gao, J.: AIM: Adversarial inference by matching priors and conditionals (2019), [https://openreview.net/forum?id=rJx\\_b3RqY7](https://openreview.net/forum?id=rJx_b3RqY7)
  36. Lim, J.H., Ye, J.C.: Geometric gan (2017)
  37. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: Proceedings of International Conference on Computer Vision (ICCV) (December 2015)
  38. Lucas, J., Tucker, G., Grosse, R.B., Norouzi, M.: Understanding posterior collapse in generative latent variable models. In: DGS@ICLR (2019)
  39. Maaløe, L., Fraccaro, M., Liévin, V., Winther, O.: Biva: A very deep hierarchy of latent variables for generative modeling (2019)
  40. Maaløe, L., Sønderby, C.K., Sønderby, S.K., Winther, O.: Auxiliary deep generative models. In: Balcan, M.F., Weinberger, K.Q. (eds.) Proceedings of The 33rd International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 48, pp. 1445–1453. PMLR, New York, New York, USA (20–22 Jun 2016), <http://proceedings.mlr.press/v48/maaloe16.html>
  41. Mao, X., Li, Q., Xie, H., Lau, R., Wang, Z., Smolley, S.: Least squares generative adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2813–2821 (10 2017). <https://doi.org/10.1109/ICCV.2017.304>
  42. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=B1QRgziT->
  43. Mukherjee, S., Asnani, H., Lin, E., Kannan, S.: Clustergan: Latent space clustering in generative adversarial networks. In: The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019. pp. 4610–4617. AAAI Press (2019).

- <https://doi.org/10.1609/aaai.v33i01.33014610>, <https://doi.org/10.1609/aaai.v33i01.33014610>
44. Ni, Y., Song, D., Zhang, X., Wu, H., Liao, L.: Cagan: Consistent adversarial training enhanced gans. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18. pp. 2588–2594. International Joint Conferences on Artificial Intelligence Organization (7 2018). <https://doi.org/10.24963/ijcai.2018/359>, <https://doi.org/10.24963/ijcai.2018/359>
  45. Nowozin, S., Cseke, B., Tomioka, R.: f-gan: Training generative neural samplers using variational divergence minimization. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) Advances in Neural Information Processing Systems 29, pp. 271–279. Curran Associates, Inc. (2016)
  46. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. In: 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings (2016), <http://arxiv.org/abs/1511.06434>
  47. Rezende, D.J., Mohamed, S., Wierstra, D.: Stochastic backpropagation and approximate inference in deep generative models. In: Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014. pp. 1278–1286 (2014), <http://jmlr.org/proceedings/papers/v32/rezende14.html>
  48. Richardson, E., Weiss, Y.: On gans and gmms. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems 31, pp. 5847–5858. Curran Associates, Inc. (2018), <http://papers.nips.cc/paper/7826-on-gans-and-gmms.pdf>
  49. Rosca, M., Lakshminarayanan, B., Mohamed, S.: Distribution matching in variational inference. CoRR [abs/1802.06847](https://arxiv.org/abs/1802.06847) (2018)
  50. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., Chen, X.: Improved techniques for training gans. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) Advances in Neural Information Processing Systems 29, pp. 2234–2242. Curran Associates, Inc. (2016), <http://papers.nips.cc/paper/6125-improved-techniques-for-training-gans.pdf>
  51. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution (2019)
  52. Springenberg, J.T.: Unsupervised and semi-supervised learning with categorical generative adversarial networks. In: 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings (2016), <http://arxiv.org/abs/1511.06390>
  53. Srivastava, A., Valkov, L., Russell, C., Gutmann, M.U., Sutton, C.: Veegan: Reducing mode collapse in gans using implicit variational learning. In: Advances in Neural Information Processing Systems. pp. 3308–3318 (2017)
  54. Tao, Y., Takagi, K., Nakata, K.: Rdec: Integrating regularization into deep embedded clustering for imbalanced datasets. In: Zhu, J., Takeuchi, I. (eds.) Proceedings of The 10th Asian Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 95, pp. 49–64. PMLR (14–16 Nov 2018), <http://proceedings.mlr.press/v95/tao18a.html>
  55. Tran, D., Ranganath, R., Blei, D.M.: Hierarchical implicit models and likelihood-free variational inference (2017)
  56. Wang, L., Schwing, A., Lazebnik, S.: Diverse and accurate image description using a variational auto-encoder with an additive gaussian encoding space. In: Guyon,

- I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 30, pp. 5756–5766. Curran Associates, Inc. (2017)
57. Wu, Y., Donahue, J., Balduzzi, D., Simonyan, K., Lillcrap, T.P.: LOGAN: latent optimisation for generative adversarial networks. *CoRR* **abs/1912.00953** (2019), <http://arxiv.org/abs/1912.00953>
  58. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747* (2017)
  59. Xie, J., Girshick, R.B., Farhadi, A.: Unsupervised deep embedding for clustering analysis. In: *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*. pp. 478–487 (2016), <http://jmlr.org/proceedings/papers/v48/xieb16.html>
  60. Xie, Q., Dai, Z., Hovy, E., Luong, M.T., Le, Q.V.: Unsupervised data augmentation for consistency training. *arXiv preprint arXiv:1904.12848* (2019)
  61. Yang, B., Fu, X., Sidiropoulos, N.D., Hong, M.: Towards k-means-friendly spaces: Simultaneous deep learning and clustering. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. pp. 3861–3870. *JMLR. org* (2017)
  62. Yang, J., Parikh, D., Batra, D.: Joint unsupervised learning of deep representations and image clusters. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5147–5156 (2016)
  63. Yang, L., Cheung, N.M., Li, J., Fang, J.: Deep clustering by gaussian mixture variational autoencoders with graph embedding. In: *The IEEE International Conference on Computer Vision (ICCV)* (October 2019)
  64. Yu, Y., Zhou, W.J.: Mixture of gans for clustering. In: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. pp. 3047–3053. *International Joint Conferences on Artificial Intelligence Organization* (7 2018). <https://doi.org/10.24963/ijcai.2018/423>, <https://doi.org/10.24963/ijcai.2018/423>
  65. Zhao, S., Song, J., Ermon, S.: The information autoencoding family: A lagrangian perspective on latent variable generative models. *arXiv preprint arXiv:1806.06514* (2018)
  66. Zong, B., Song, Q., Min, M.R., Cheng, W., Lumezanu, C., Cho, D., Chen, H.: Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In: *International Conference on Learning Representations* (2018), <https://openreview.net/forum?id=BJJLHbb0->