

# Top-Down Deep Clustering with Multi-Generator GANs

Daniel P. M. de Mello<sup>1</sup>, Renato M. Assunção<sup>2,1</sup>, Fabricio Murai<sup>1</sup>

<sup>1</sup> Universidade Federal de Minas Gerais,  
<sup>2</sup> Esri Inc.

dani-dmello@hotmail.com, rassuncao@esri.com, murai@dcc.ufmg.br

## Abstract

Deep clustering (DC) leverages the representation power of deep architectures to learn embedding spaces that are optimal for cluster analysis. This approach filters out low-level information irrelevant for clustering and has proven remarkably successful for high dimensional data spaces. Some DC methods employ Generative Adversarial Networks (GANs), motivated by the powerful latent representations these models are able to learn implicitly. In this work, we propose HC-MGAN, a new technique based on GANs with multiple generators (MGANs), which have not been explored for clustering. Our method is inspired by the observation that each generator of a MGAN tends to generate data that correlates with a sub-region of the real data distribution. We use this clustered generation to train a classifier for inferring from which generator a given image came from, thus providing a semantically meaningful clustering for the real distribution. Additionally, we design our method so that it is performed in a top-down hierarchical clustering tree, thus proposing the first hierarchical DC method, to the best of our knowledge. We conduct several experiments to evaluate the proposed method against recent DC methods, obtaining competitive results. Last, we perform an exploratory analysis of the hierarchical clustering tree that highlights how accurately it organizes the data in a hierarchy of semantically coherent patterns.

## Introduction

Cluster analysis is a fundamental problem in unsupervised learning, with a wide range of applications, especially in computer vision (Shi and Malik 2000; Achanta and Susstrunk 2017; Joulin, Bach, and Ponce 2010; Liu et al. 2018). Its goal is to assign similar points of the data space to the same cluster, while ensuring that dissimilar points are placed in different clusters. One of the main challenges in this approach is to quantify the similarity between objects. For low-dimensional data spaces, similarity might be straightforwardly defined as the minimization of some geometric distance (*e.g.* euclidean distance, squared euclidean distance, Manhattan distance). On the other hand, choosing the distance metric becomes unfeasible for high dimensional data distributions. Images are a clear example of this problem, since any distance metric based on raw pixel spaces

is subject to all sorts of low-level noisy disturbances irrelevant for determining the similarity and suffer from a lack of translation or rotation invariance. This motivates the need for some dimensionality reduction technique, by which the fundamental relationships between objects projected onto the resulting embedding space become more consistent with geometrical distances.

In recent years deep clustering techniques have spearheaded the dimensionality reduction approach to clustering by employing highly non-linear latent representations learned by deep learning models (Krizhevsky, Sutskever, and Hinton 2012). Considering the unsupervised nature of cluster analysis, the models that naturally arise as candidates for deep clustering are unsupervised deep generative models, since these must learn highly abstract representations of the data as a requirement for realistic and diverse generated samples. One of such models are the Generative Adversarial Networks (GAN) (Goodfellow et al. 2014), whose extremely realistic results in image generation, semantic interpolation and interpretability in the latent space, are evidence of their capacity of learning a powerful latent representation that captures the essential components of the data distribution. Nonetheless, very few works have proposed GAN architectures designed for clustering. Some of these works are (Mukherjee et al. 2019; Chen et al. 2016), where the authors showed that, by manipulating the generator’s architecture in a specific way, it is possible to control the class of the training data to which a generated sample belongs, even when classes labels are not available during the training.

Some recent works employed a GAN architecture with multiple generators (Ghosh et al. 2018; Hoang et al. 2018; Zhang et al. 2018) to achieve greater diversity in image generation, as well as an alternative way of stabilizing the training. In these works, the authors have observed that each generator tends to specialize in generating examples belonging to a specific class of the dataset. This suggests the organic emergence of clusters in the generators’ representation of the data and it is one of the main motivations of our work. The rationale is that clustering would be possible by employing a classifier in charge of distinguishing between the generators, and this classifier could later be applied to the real dataset in order to classify real examples without the use of labels. Additionally, the problem of setting the number of generators to be used for representing the classes of the training set is

not addressed in these works, which could be an issue in a real clustering task, where the number of clusters is assumed to be unknown.

In this work, we propose Hierarchical Clustering MGAN (HC-MGAN), a method that leverages the multi-generator GAN for the clustering task. We employ multiple generators, each of them specializing in representing a particular cluster of the training distribution. This should lead to a stronger representation capacity and with more meaningful clusters than what a single generator covering multiple clusters can provide. MGANs have not been used in the previous works exploring GANs for clustering. Additionally, we design our method so that it performs the clustering of the training data in a top-down hierarchical way, creating new generators as divisions of subsequent clusters become necessary, *i.e.*, it permits the user to control different clustering granularity levels according to the task at hand. The main contributions of this work are as follows:

- We propose HC-MGAN, a novel deep clustering method that employs GANs with multiple generators.
- We design HC-MGAN as a top-down hierarchical clustering tree introducing the first hierarchical deep clustering algorithm. Hierarchical clustering allows the user to control different levels of cluster granularity and favors interpretability by taxonomically organizing the clusters.
- We conduct experiments with three image datasets used in other deep clustering works as benchmarks, namely, MNIST, Fashion MNIST and Stanford Online Products. We obtain competitive results against recent deep clustering methods, all of which are horizontal and therefore lack the advantages of the hierarchical approach.
- We explore the clustering pattern obtained throughout the tree, displaying how HC-MGAN is able to organize the data in a semantically coherent hierarchy of clusters.

## Related Work

We review some recent work in topics related to Deep Clustering and GANs. Autoencoders (AEs) have been the most prevalent choice in the deep clustering literature (Xie, Girshick, and Farhadi 2016; Yang et al. 2017, 2019; Zhang et al. 2019b), where the clustering objective is usually optimized on the feature representation  $\mathbf{Z}$  resultant of the mapping  $E: \mathcal{X} \rightarrow \mathcal{Z}$  learned by the encoder component in AEs for a data space  $\mathcal{X}$  and a feature space  $\mathcal{Z}$ . Deep Embedded Clustering (DEC) (Xie, Girshick, and Farhadi 2016) pioneered this approach, obtaining state-of-the-art clustering results. It works by pretraining an AE with a standard reconstruction loss function, and then optimizing it with a regularizer based on the clustering assignments modeled by a target student-distribution, having the cluster centers iteratively updated. DEC’s results were surpassed by (Yang et al. 2017), which converted DEC’s objective function into an alternated joint optimization with K-Means loss, thus obtaining a clusterization more suited for K-Means.

The use of GANs for clustering tasks has been influenced by InfoGAN (Chen et al. 2016), a type of GAN whose latent variable consists of, besides the usual multidimensional variable  $z$ , a set  $c$  of one-dimensional variables  $c_1, c_2, \dots, c_N$

that are expected to unsupervisedly capture semantic information in a disentangled manner (*i.e.*, with each variable encoding isolated interpretable features of the real data). For obtaining this, the authors of InfoGAN proposed an additional term in the generator’s loss function that maximized the mutual information  $I(c; G(z, c))$  between a generated image  $G(z, c)$  and the latent variable  $c$  that originated it. The variables  $c$  could be chosen to represent both categorical and continuous features. ClusterGAN (Mukherjee et al. 2019) is an architecture appropriate for clustering tasks based on InfoGAN. The generator learns to generate a certain class of the real distribution correlated with a given one-hot format for the latent variables  $c$ . To obtain this, they proposed to use an inference encoder network capable of performing the mapping  $E: \mathcal{X} \rightarrow \mathcal{Z}$ , which is the inverse of the generator’s mapping and similar to an encoder’s mapping for an autoencoder architecture. After the training, the encoder can be employed to classify real data samples according to the latent variable to which it is mostly correlated, thus providing the clustering. Two key differences of our work is that (i) we use a separate generator (not a discrete latent variable) to encode a cluster, enabling our method to discover clusters with more representation capacity, and that (ii) we obtain hierarchical clusters, unlike the previous methods.

Kundu et al. (2019) propose the GAN-Tree framework, which slightly resembles our approach, since it also involves a hierarchical structure of independent nodes containing GANs capable of generating samples related to different levels of a similarity hierarchy. There are several differences, however. The most important is that the main motivation for GAN-Tree was a framework capable of addressing the trade-off between quality and diversity when generating samples from multi-modal data. The authors claimed that their approach could be readily adapted for clustering tasks, but no definitive experiments with clustering benchmarks were provided. Other important difference lies in their splitting procedure, which was performed with a latent  $\hat{z}$  inferred by an encoder  $E$  for a sample image  $\mathbf{x}$ , that is,  $\hat{z} = E(\mathbf{x})$ . For each node of the tree, they decompose their prior for  $\hat{z}$  into a mixture of two Gaussians with shifted means. They determine the prior Gaussian component to which  $\hat{z}$  is more likely related, and then train the encoder to maximize the likelihood to this prior. For a clustering task, this approach would heavily rely on the assumption that the inference made by  $E$ , as well as the cluster encoding with the decomposed Gaussians in  $\mathcal{Z}$ , will be sufficient to capture semantically meaningful clustering patterns. The split in our approach, on the other hand, is directly embedded into the GAN training, with each generator automatically learning to represent each cluster. Therefore, in our work the clustering semantic quality is directly tied to the GAN’s well known representation learning capacity, and, in particular, to the tendency of different generators in MGANs to cover different areas of the training distribution with high semantic discrepancy.

## Proposed Method: HC-MGAN

First, we provide a bird’s eye view of the proposed method. Then we describe how its two phases – Raw Split and Refinement – work in detail.




Figure 1: Hierarchical clustering tree overview on MNIST. Nodes are represented by membership vectors  $\mathbf{s}_k$ . Grids show images sampled from each distribution  $\mathcal{P}_k$ . Each split divides the probability masses in  $\mathbf{s}_k$  into two new nodes.

#### Algorithm 1: Split

**Input:**  $X_{\text{Data}}, \mathbf{s}_k$

- 1:  $\mathbf{s}_l^{(0)}, \mathbf{s}_m^{(0)} \leftarrow \text{raw\_split}(X_{\text{Data}}, \mathbf{s}_k)$
- 2: **for**  $t = 1, \dots, T$  **do**
- 3:    $\mathbf{s}_l^{(t)}, \mathbf{s}_m^{(t)} \leftarrow \text{refinement}(X_{\text{Data}}, \mathbf{s}_l^{(t-1)}, \mathbf{s}_m^{(t-1)})$
- 4: **return**  $\mathbf{s}_l = \mathbf{s}_l^{(T)}, \mathbf{s}_m = \mathbf{s}_m^{(T)}$

### Overview of the Hierarchical Scheme

For a given a collection of  $N$  training examples  $X_{\text{Data}} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ , our method constructs a binary tree of hierarchical clusters, iteratively creating nodes, from top to bottom. The  $k$ -th created node is represented by a vector  $\mathbf{s}_k \in \mathbb{R}^N$ , referred to as membership vector, where each  $s_{k,i} = p(Z_i = k | \mathbf{x}_i, \theta)$  measures the probability of example  $i$  belonging to cluster  $k$  given  $\mathbf{x}_i$  and our model’s parameters  $\theta$ , *i.e.*, a soft clustering approach. The initial  $\mathbf{s}_0$  consists of an all-ones vector. Figure 1 depicts the initial growth of the tree with the MNIST dataset, for the first 7 nodes. The tree grows via a split: a soft clustering operation that takes as input a node  $\mathbf{s}_k$ , and divides its probability masses into two new nodes represented by membership vectors  $\mathbf{s}_l$  and  $\mathbf{s}_m$ , with  $\mathbf{s}_l + \mathbf{s}_m = \mathbf{s}_k$ . We decide which leaf to split next by taking the node  $\mathbf{s}_k$  with the largest total mass.

The split mechanism consists of two phases: (i) **raw split** and (ii) **refinement phase**, as shown in Algorithm 1. Before we obtain the final  $\mathbf{s}_l$  and  $\mathbf{s}_m$  vectors, vector  $\mathbf{s}_k$  undergoes a raw split, which outputs the initial membership vectors  $\mathbf{s}_l^{(0)}$  and  $\mathbf{s}_m^{(0)}$ . In most cases,  $\mathbf{s}_l^{(0)}$  and  $\mathbf{s}_m^{(0)}$  are just rough estimates of how to split  $\mathbf{s}_k$  in two clusters. Hence, we use the refinement phase to get them progressively closer to what we expect the ideal soft clustering assignment to be. In Algorithm 1, we can see how the refinement transforms  $\mathbf{s}_l^{(0)}$  and  $\mathbf{s}_m^{(0)}$  into two new membership vectors  $\mathbf{s}_l^{(1)}$  and  $\mathbf{s}_m^{(1)}$ . This process is repeated for  $T$  refinement operations to yield the final result  $\mathbf{s}_l = \mathbf{s}_l^{(T)}$  and  $\mathbf{s}_m = \mathbf{s}_m^{(T)}$ . Note that  $\mathbf{s}_m^{(t)} + \mathbf{s}_l^{(t)} = \mathbf{s}_k$  for every  $t$ .

Figure 2 shows the progressive improvement resulting from the refinement phase with a grid of 25 MNIST samples, for an  $\mathbf{s}_k$  used as a running example. Samples of each membership vector are shown below its label and color-coded (gray scale) according to their probability mass. In this example, 3’s and 5’s are the classes mostly associated with  $\mathbf{s}_k$ . We expect the final split result ( $\mathbf{s}_l$  and  $\mathbf{s}_m$ ) to be a separa-

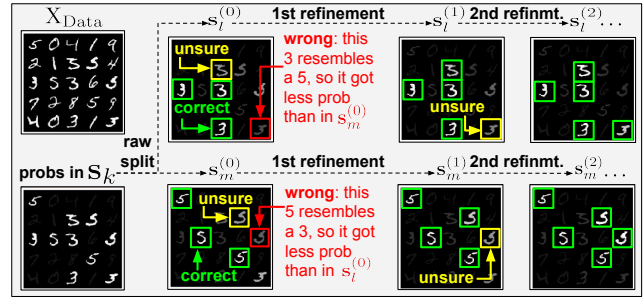


Figure 2: (Best viewed in color) Split mechanism applied to running example  $\mathbf{s}_k$  from MNIST. Grids show samples color-coded (gray scale) according to their probability mass in each membership vector (white: 100%).  $\mathbf{s}_k$  is mostly associated with 3’s and 5’s. Raw split roughly separates both classes. Refinement iterations enhance the separation.

tion of the probability mass of 3’s and 5’s. After the raw split of  $\mathbf{s}_k$ , the probability mass of 3’s and 5’s is roughly divided between  $\mathbf{s}_l^{(0)}$  and  $\mathbf{s}_m^{(0)}$ , but the most ambiguous examples received low probability mass in the membership vector mostly associated to its class. After the first refinement, we observe that  $\mathbf{s}_l^{(1)}$  and  $\mathbf{s}_m^{(1)}$  provide a better separation of 3’s and 5’s. After another iteration, most of the probability mass of the 3’s (5’s) samples is in  $\mathbf{s}_l^{(2)}$  ( $\mathbf{s}_m^{(2)}$ ).

### Raw Split

For the raw split, we use a two-generator MGAN architecture, which we adapt for binary clustering by leveraging the fact that each generator learns to specialize in generating samples from one sub-region of the real data distribution, typically correlated with a specific set of classes of the data.

Figure 3 (Top) depicts the training of this MGAN for our running example. The MGAN components are: generators  $G\alpha_k, G\beta_k$ , discriminator  $D_k$  and classifier  $C_k$ . We need each generator to specialize in sub-regions of  $\mathbf{s}_k$ . Hence, real data samples used for training them must reflect the sample distribution given by  $\mathbf{s}_k$  (in the example, these are mostly 3’s and 5’s). We define distribution  $\mathcal{P}_{\mathbf{s}_k}$  over the real examples by (sum-to-one) normalizing  $\mathbf{s}_k$ , and use it to sample the training batches. We train the MGAN with the usual adversarial game between generators  $G\alpha_k, G\beta_k$  and the discriminator  $D_k$ , while  $C_k$  is trained to distinguish between generators. After a few epochs, we observe that one generator is generating mostly 3’s while the other, mostly 5’s. An important detail is that we also train the generators to minimize  $C_k$ ’s classification loss, increasing the incentive for  $G\alpha_k$  and  $G\beta_k$  to generate samples from different sub-regions. Moreover, we share some parameters between  $D_k$  and  $C_k$ , since it is more desirable for  $C_k$  to perform its classification in a higher-level feature space, such as the one learned by  $D_k$ .

After training the MGAN for enough epochs, we can cluster the data as shown in Figure 3 (Bottom). We use  $C_k$  to perform the clustering on the real data according to the two generated distributions it learned to identify, and which

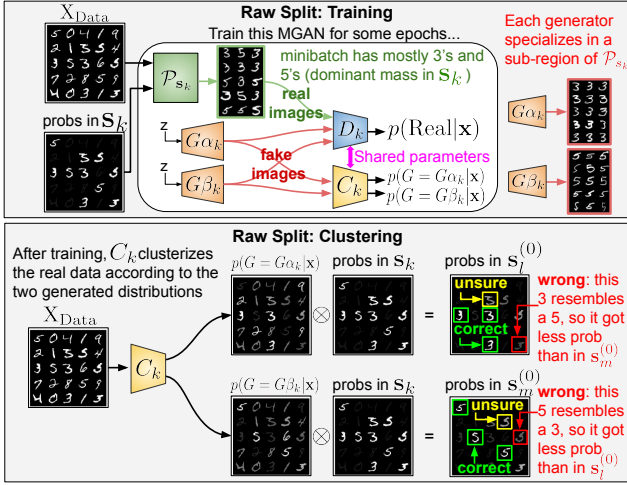


Figure 3: (Top) Training the Raw Split Components: generators  $G_{\alpha_k}$ ,  $G_{\beta_k}$ , discriminator  $D_k$  and classifier  $C_k$ .  $\mathcal{P}_{s_k}$  draws real samples in proportion to their mass in  $s_k$ . (Bottom) Clustering the dataset with the raw split classifier.

we expect to correlate with different classes of the dataset. In this example, the first generated distribution resembled 3's, while the other resembled 5's. Hence, we expect  $C_k$  to mostly assign real 3's probability mass to the first soft cluster and real 5's probability mass to the second. As seen in Figure 3 (Bottom),  $C_k$  does manage to roughly assign the 3's (resp. 5's) mass to the cluster related to  $G_{\alpha_k}$ 's (resp.  $G_{\beta_k}$ ).

Note that  $C_k$  must be used to perform inference for all examples of every class in the dataset, regardless if they were present in the generated distributions, with  $p(G = G_{\alpha_k}|\mathbf{x}) + p(G = G_{\beta_k}|\mathbf{x}) = 1$ . While  $C_k$  can be seen as classifying examples conditioned on them belonging to node's  $s_k$  subtree, membership vectors actually correspond to estimates for unconditional probabilities. To enforce the condition that, for each example, the sum of its probability masses in the membership vectors of  $s_k$ 's children equals vector  $s_k$ , we multiply the probabilities of each example predicted by  $C_k$  by its probability in  $s_k$ .

Finally, by noting that some samples in the running example were not well separated in  $s_l^{(0)}$  and  $s_m^{(0)}$  by classifier  $C_k$ , we conclude that the two distributions obtained from a raw split may not be sufficiently diverse or have enough quality to account for the entire set of 3's and 5's the MGAN had access to. This underlines the need for the refinement phase, which will be covered in the next section.

We now formalize the MGAN game in the raw split phase. Dropping the subscript  $k$  to avoid clutter, we define the objective function of the two-generator MGAN a sum of two cost functions  $\mathcal{L}_{adv}$  and  $\mathcal{L}_{cls}$ ,

$$\min_{G_{\alpha}, G_{\beta}, C} \max_D \mathcal{L} = \mathcal{L}_{adv}(G_{\alpha}, G_{\beta}, D) + \lambda \mathcal{L}_{cls}(G_{\alpha}, G_{\beta}, C), \quad (1)$$

where  $\mathcal{L}_{adv}$  is the cost for the adversarial minimax game

## Algorithm 2: Raw Split

**Input:**  $X_{Data}, s_k$

- 1: Creates Components  $G_{\alpha_k}, G_{\beta_k}, C_k, D_k$
- 2: **for** training\_iterations **do**
- 3: sample  $\mathbf{x}, \hat{\mathbf{x}}_{\alpha}, \hat{\mathbf{x}}_{\beta}$  from  $\mathcal{P}_{s_k}, \mathcal{P}_{G_{\alpha_k}}, \mathcal{P}_{G_{\beta_k}}$
- 4: Get  $\mathcal{L}_{D_k}^{(real)}$  using  $D_k$  criterion on  $\mathbf{x}$  with real labels
- 5: Get  $\mathcal{L}_{D_k}^{(fake)}$  using  $D_k$  criterion on  $\hat{\mathbf{x}}_{\alpha}, \hat{\mathbf{x}}_{\beta}$  w/ fake labels
- 6: Update  $\theta_{D_k}$  with Adam( $\nabla_{\theta_{D_k}}(\mathcal{L}_{D_k}^{(real)} + \mathcal{L}_{D_k}^{(fake)})$ )
- 7: Get  $\mathcal{L}_{C_k}$  using  $C_k$  criterion on  $\hat{\mathbf{x}}_{\alpha}, \hat{\mathbf{x}}_{\beta}$ , with categorical labels for  $G_{\alpha}, G_{\beta}$
- 8: Update  $\theta_{C_k}$  with Adam( $\nabla_{\theta_{C_k}}(\mathcal{L}_{C_k})$ )
- 9: Get  $\mathcal{L}_{G_k}^{(disc)}$  using  $D_k$  criterion on  $\hat{\mathbf{x}}_{\alpha}, \hat{\mathbf{x}}_{\beta}$  with real labels
- 10: Get  $\mathcal{L}_{G_k}^{(clasf)}$  using  $C_k$  criterion on  $\hat{\mathbf{x}}_{\alpha}, \hat{\mathbf{x}}_{\beta}$  with categorical labels for  $G_{\alpha}, G_{\beta}$
- 11: Update  $\theta_{G_{\alpha_k}, G_{\beta_k}}$  w/ Adam( $\nabla_{\theta_{G_k}}(\mathcal{L}_{G_k}^{(disc)} + \lambda \mathcal{L}_{G_k}^{(clasf)})$ )
- 12: **for**  $\mathbf{x}_i$  in  $X_{Data}$  **do**
- 13:  $s_{l,i} \leftarrow C_k^{(\alpha, out)}(\mathbf{x}_i) \cdot s_{k,i}, s_{m,i} \leftarrow C_k^{(\beta, out)}(\mathbf{x}_i) \cdot s_{k,i}$
- 14: **return**  $s_l, s_m$

between generators and discriminator, given by

$$\mathcal{L}_{adv}(G_{\alpha}, G_{\beta}, D) = \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_s} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_{G_{\alpha}}} [\log(1 - D(\mathbf{x}))] + \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_{G_{\beta}}} [\log(1 - D(\mathbf{x}))] \quad (2)$$

and  $\mathcal{L}_{cls}$  is the classification cost that is minimized by both generators and classifier, given by

$$\mathcal{L}_{cls}(G_{\alpha}, G_{\beta}, C) = \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_{G_{\alpha}}} [\log(C(\mathbf{x}))] + \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_{G_{\beta}}} [\log(C(\mathbf{x}))]. \quad (3)$$

Note that we multiply  $\mathcal{L}_{cls}$  by a regularization parameter  $\lambda$  to weight its impact on the total cost.

Algorithm 2 lists the steps involved in training the MGAN for the raw split and the clustering performed with  $C_k$ .

## Refinement

The refinement phase comes right after the raw split. During this phase, some of the probability mass in the two membership vectors obtained by the raw split,  $s_l^{(0)}$  and  $s_m^{(0)}$ , is exchanged so as to iteratively improve the clustering quality (see Figure 2). Each iteration is referred to as a refinement sub-block. Without loss of generality, consider the first refinement sub-block, whose components are depicted in Figure 4 (Top), for our running example. The components are divided in two "refinement groups",  $l$  and  $m$  (each formed by a generator  $G_j$ , a discriminator  $D_j$  and a classifier  $C_j$ ,  $j \in \{l, m\}$ ). Each group  $j$  takes  $s_j^{(0)}$  as input, and has its own independent GAN game occurring between  $G_j$  and  $D_j$ .

This scheme with two separated GANs is designed to obtain a more focused generative representation of each sub-region of  $s_k$  than we were able to obtain at the raw split phase using a single MGAN's discriminator to learn to discriminate the entire region described by  $s_k$ . *By providing a more focused view of one sub-region to one discriminator, it encounters less variance among the real examples it receives, and thus its discriminative task becomes easier. We expect its adversarial generator's response to be a more diverse and convincing generation of examples associated with that particular sub-region.*

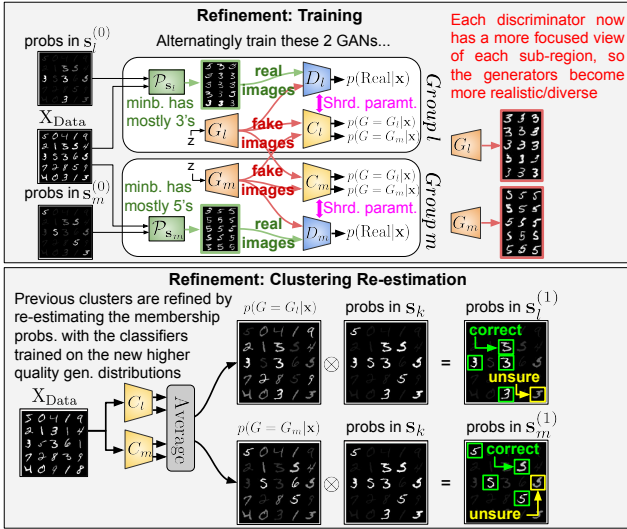


Figure 4: (Top) Refinement Training: generators  $G_l, G_m$ , discriminators  $D_l, D_m$  and classifiers  $C_l, C_m$ . (Bottom) Clustering the dataset with the refinement classifiers.

As shown in Figure 4 (Top), each GAN in groups  $j \in \{l, m\}$  draws real samples from its corresponding distribution  $\mathcal{P}_{s_j}$  over  $X_{\text{Data}}$ , which is equal to the (sum-to-one) normalized vector  $s_j$  (analogously to  $\mathcal{P}_{s_k}$ ). As a result, the minibatches passed to each group’s GAN reflect the probability mass in their respective membership vectors, *e.g.*: in our running example,  $\mathcal{P}_{s_l}$  draws mostly 3’s, since 3’s have more probability mass in  $s_l^{(0)}$ , but it might eventually draw some 5’s as well, since there’s still some mass for 5’s in  $s_l^{(0)}$ .

The role of classifiers  $C_l$  and  $C_m$  in these two separated GAN games is similar to the single classifier  $C_k$  used in the raw split phase: learn to distinguish samples from  $G_l$  and  $G_m$ , thereby providing a way to cluster the real data. However, instead of using a single classifier (as in the raw split), we found that using two separate classifiers which share parameters with the respective discriminators forces the classification to occur in a higher-level feature space, achieving better clustering. Furthermore, we also train  $G_l$  and  $G_m$  to minimize both classifiers’ losses, thus providing incentive for each generator to generate samples more strongly correlated with each sub-region (we did something similar for the 2 generators in the MGAN of the raw split).

After training the two refinement groups alternately for enough epochs, we perform a **clustering re-estimation**, as shown in Figure 4 (Bottom). This is similar to the way  $C_k$  was used to cluster the real data in the raw split (Figure 3 (Bottom)). Now, both classifiers estimate the probability that a sample  $x_i \in X_{\text{Data}}$  came from  $G_l$  (or its complement,  $G_m$ ), hereby denoted by  $C_j^{(l.out)}(x_i)$ , for  $j \in \{l, m\}$ . We take the average probability  $(C_l^{(l.out)}(x_i) + C_m^{(l.out)}(x_i))/2$  as the proportion of the mass associated with  $x_i$  in  $s_k$  that should go into  $s_l^{(1)}$  (complement to  $s_m^{(1)}$ ). By using this “two-classifier ensemble”, we aim to increase the quality of the

clustering, since the generators’ distributions are expected to be more representative of each sub-region of the dataset than the two generated distributions obtained during the raw split. In the running example,  $G_l$ ’s ( $G_m$ ’s) distributions resembled 3’s (5’s), so the classifiers assign more of the 3’s (5’s) probability mass to the cluster associated with  $G_l$  ( $G_m$ ’s).

For the subsequent refinement, we expect that using  $s_l^{(1)}$  and  $s_m^{(1)}$  to train new refinement groups  $l$  and  $m$  can yield generated distributions that are even more representative of the sub-regions encoded by these membership vectors, providing, in turn, even more information for the classifiers to perform the clustering and to obtain improved membership vectors  $s_l^{(2)}$  and  $s_m^{(2)}$  (recall Figure 2). Therefore, by repeating the process over  $T$  refinements, it is expected that the initial sub-regions captured by  $s_k$  are increasingly more associated with either of the refinement groups.

We now formalize the two simultaneous GAN games for the refinement training. From the perspective of refinement group  $l$ , the training can be defined as an optimization of a sum of two cost functions  $\mathcal{L}_{adv}$  and  $\mathcal{L}_{cls}$ ,

$$\min_{G_l, C_l} \max_{D_l} \mathcal{L}(G_l, D_l, C_l) = \mathcal{L}_{adv}(G_l, D_l) + \lambda \mathcal{L}_{cls}(G_l, C_l), \quad (4)$$

where  $\mathcal{L}_{adv}$  describes the cost function for the adversarial minimax game between generator  $G_l$  and discriminator  $D_l$ , that only involves group  $l$  components, and is given by

$$\mathcal{L}_{adv}(G_l, D_l) = \mathbb{E}_{x \sim \mathcal{P}_{data}} [\log D_l(x)] + \mathbb{E}_{x \sim \mathcal{P}_{G_l}} [\log(1 - D_l(x))] \quad (5)$$

and  $\mathcal{L}_{cls}$  is classification cost that is minimized with respect to  $G_l$ ’s parameters and  $C_l$ ’s parameters, but also involves  $G_m$  and  $C_m$  for computing the cost, given by

$$\mathcal{L}_{cls}(G_l, C_l) = \mathbb{E}_{x \sim \mathcal{P}_{G_l}} [\log C_l(x)] + \mathbb{E}_{x \sim \mathcal{P}_{G_l}} [\log C_m(x)] + \mathbb{E}_{x \sim \mathcal{P}_{G_m}} [\log C_l(x)]. \quad (6)$$

We multiply  $\mathcal{L}_{cls}$  by a regularization parameter  $\lambda$  to weight its impact on the total cost. The corresponding equations for refinement group  $m$  follow analogously.

Algorithms 3 and 4 respectively list the steps involved in the external training loop that coordinates the alternated training of refinement groups  $l$  and  $m$ , and in the function called by Algorithm 3 that performs an update for the components of a given refinement group isolatedly.

## Experiments

As unsupervised tasks forbid hyperparameter tuning, we used only slightly different tunings for each dataset, none of which required labeled supervision, merely aiming to stabilize the generators and avoid overfitting the classification of generators. Code available at [github.com/dmdmello/HC-MGAN](https://github.com/dmdmello/HC-MGAN) and supplementary/implementation details at [arxiv.org/abs/2112.03398](https://arxiv.org/abs/2112.03398).

We consider three datasets: MNIST (LeCun et al. 1998), Fashion MNIST (FMNIST) (Xiao, Rasul, and Vollgraf 2017) and Stanford Online Products (SOP) (Oh Song et al. 2016). Following a common practice in DC works, we used all available images for each dataset. **MNIST**: This dataset consists of grayscale images of hand-written digits with 28x28 resolution, with 10 classes and approximately 7k im-

---

**Algorithm 3: Refinement**

---

**Input:**  $X_{\text{Data}}, \mathbf{s}_l^{(t)}, \mathbf{s}_m^{(t)}$

- 1: Creates  $G_l, D_l, C_l$  for group  $l$  and  $G_m, D_m, C_m$  for group  $m$
- 2: **for**  $T$  iterations **do**
- 3: sample  $\mathbf{x}_l, \mathbf{x}_m, \hat{\mathbf{x}}_l, \hat{\mathbf{x}}_m$  from  $\mathcal{P}_{\mathbf{s}_l^{(t)}}, \mathcal{P}_{\mathbf{s}_m^{(t)}}, \mathcal{P}_{G_l}, \mathcal{P}_{G_m}$
- 4: TrainRefinGroup( $\mathcal{G}_{\text{int}} = \{G_l, D_l, C_l, \mathbf{x}_l, \hat{\mathbf{x}}_l\}, \mathcal{G}_{\text{ext}} = \{C_m, \hat{\mathbf{x}}_m\}$ ) {trains  $l$  with needed external data/components from  $m$ }
- 5: TrainRefinGroup( $\mathcal{G}_{\text{int}} = \{G_m, D_m, C_m, \mathbf{x}_m, \hat{\mathbf{x}}_m\}, \mathcal{G}_{\text{ext}} = \{C_l, \hat{\mathbf{x}}_l\}$ ) {trains  $m$  with needed external data/components from  $l$ }
- 6: **for**  $\mathbf{x}_i$  in  $X_{\text{Data}}$  **do**
- 7:  $s_{i,l}^{(t+1)} \leftarrow (C_l^{(l.out)}(\mathbf{x}_i) + C_m^{(l.out)}(\mathbf{x}_i)) \cdot (s_{i,l}^{(t)} + s_{i,m}^{(t)})/2$
- 8:  $s_{i,m}^{(t+1)} \leftarrow (C_l^{(m.out)}(\mathbf{x}_i) + C_m^{(m.out)}(\mathbf{x}_i)) \cdot (s_{i,l}^{(t)} + s_{i,m}^{(t)})/2$
- 9: **return**  $\mathbf{s}_l^{(t+1)}, \mathbf{s}_m^{(t+1)}$

---

---

**Algorithm 4: TrainRefinGroup**

---

**Input:**  $\mathcal{G}_{\text{int}} = \{G_{\text{int}}, D_{\text{int}}, C_{\text{int}}, \mathbf{x}, \hat{\mathbf{x}}_{\text{int}}\}, \mathcal{G}_{\text{ext}} = \{C_{\text{ext}}, \hat{\mathbf{x}}_{\text{ext}}\}$

# $\mathcal{G}_{\text{int}}$  receives internal data/components from the current refinement group being trained,  $\mathcal{G}_{\text{ext}}$  receives external data/components from the neighbor refin. group needed to train the current group

- 1: Get  $\mathcal{L}_{D_{\text{int}}}^{(real)}$  using  $D_{\text{int}}$  criterion on  $\mathbf{x}$  with real labels
- 2: Get  $\mathcal{L}_{D_{\text{int}}}^{(fake)}$  using  $D_{\text{int}}$  criterion on  $\hat{\mathbf{x}}_{\text{int}}$  with fake labels
- 3: Updates  $\theta_{D_{\text{int}}}$  with Adam( $\nabla_{\theta_{D_{\text{int}}}}(\mathcal{L}_{D_{\text{int}}}^{(real)} + \mathcal{L}_{D_{\text{int}}}^{(fake)})$ )
- 4: Get  $\mathcal{L}_{C_{\text{int}}}$  using  $C_{\text{int}}$  criterion on  $\hat{\mathbf{x}}_{\text{int}}, \hat{\mathbf{x}}_{\text{ext}}$ , with categorical labels for internal and external generated data
- 5: Updates  $\theta_{C_{\text{int}}}$  with Adam( $\nabla_{\theta_{C_{\text{int}}}}(\mathcal{L}_{C_{\text{int}}})$ )
- 6: Get  $\mathcal{L}_{G_{\text{int}}}^{(disc)}$  using  $D_{\text{int}}$  criterion on  $\hat{\mathbf{x}}_{\text{int}}$  with real labels
- 7: Get  $\mathcal{L}_{G_{\text{int}}}^{(clasf)}$  using  $C_{\text{int}}$  criterion on  $\hat{\mathbf{x}}_{\text{int}}, \hat{\mathbf{x}}_{\text{ext}}$  with categorical labels for internal and external generated data
- 8: Updates  $\theta_{G_{\text{int}}}$  with Adam( $\nabla_{\theta_{G_{\text{int}}}}(\mathcal{L}_{G_{\text{int}}}^{(disc)} + \lambda \mathcal{L}_{G_{\text{int}}}^{(clasf)})$ )

---

ages available for each class. **FMNIST:** This dataset consists of gray scale pictures of clothing-related objects with 28x28 resolution, with 10 classes and exactly 7k images available for each class. **SOP:** This dataset consists of color pictures of products with varying resolution sizes, with 12 classes, and a varied number of examples per class, roughly ranging from 6k examples to 13k. SOP, in particular, is designed for supervised tasks, and is very hard for clustering due to high intra-class variance. We follow the practice adopted for SOP in (Zhang et al. 2019b), *i.e.*, we convert the images to grayscale, resize them to 32x32, and remove the classes “kettle” and “lamp”.

**Evaluation Metrics** We consider two of the most common clustering metrics for benchmarks with each dataset: *clustering accuracy* (ACC) and *normalized mutual information* (NMI). As usual, these metrics are computed on the results obtained when setting the number of clusters  $C$  to the number of classes in the data. For a direct comparison, we let HC-MGAN’s tree grow until it reaches  $C$  leaves. Additionally, we convert the final soft clustering outputs to hard assignments by attributing each example to the highest probability group, and then compute the evaluation metrics.

**Baselines and SOTA methods** We present the baselines and state-of-the-art methods used in the comparison. We consider five groups of methods: (i) classical, non Deep Learning (DL)-based: K-Means (MacQueen 1967), SC (Zelnik-Manor and Perona 2004), AC (Gowda and Krishna 1978), NMF (Cai et al. 2009); (ii) Varied DC: DEC (Xie, Girshick, and Farhadi 2016), DCN (Yang et al. 2017), JULE (Yang, Parikh, and Batra 2016), VaDE (Jiang et al. 2017), DEPICT (Ghasedi Dizaji et al. 2017), SpectralNET, DAC (Chang et al. 2017), (Shaham et al. 2018), DualAE (Yang et al. 2019); (iii) Subspace Clustering (either DL-based or not): NCSC (Zhang et al. 2019b), SSC (Elhamifar and Vidal 2013), LLR (Liu et al. 2013), KSSC (Patel and Vidal 2014), DSC-Net (Ji et al. 2017),  $k$ -SCN (Zhang et al. 2019a); (iv) GAN-based DC: ClusterGAN (Mukherjee et al. 2019), InfoGAN (Chen et al. 2016), DLS-Clustering (Ding and Luo 2019); (v) DC w/ data augmentation: (Zhang et al. 2019b), IIC (Ji, Henriques, and Vedaldi 2019), DCCM (Wu et al. 2019) and DCCS (Zhao et al. 2020). *None of the DC methods we found in the literature are hierarchical.* Most results are transcribed from either (Zhao et al. 2020), (Zhang et al. 2019b) or (Mukherjee et al. 2019).

**Results** Table 1 shows the performance comparison between traditional baselines, state-of-the-art DC methods without and with data augmentation, and our method. In order to check the effectiveness of the refinements, we also display results obtained only with raw splits.

**MNIST** Our method outperforms the traditional baselines by a large margin. In terms of ACC, it is not among the top 5 presented methods, but it performs reasonably close to them, even outperforming, in either NMI or ACC, some DC methods like ClusterGAN, InfoGAN, DEC, DCN. This result was obtained by employing the same architecture we used for FMNIST, which might be of excessive capacity for MNIST and even harm the clustering result by making it trivial for a single generator to represent all the data, instead of a cluster of similar data points, during the raw split. This dataset was the one for which the refinements showed the greatest improvement over the raw split only experiment.

**FMNIST** Only DCCS is able to surpass our method’s ACC and NMI performance. We emphasize that using data augmentation in DCCS causes a significant improvement, since selecting the right type of augmentations for a specific dataset can reduce much of the intra-class variance. However, augmentation with GANs is challenging (Karras et al. 2020), so we leave this for future work. The refinements still had a positive impact over the raw split only experiment.

**SOP** The results for existing methods were transcribed from the NCSC work (Zhang et al. 2019b). There the authors state that they handpicked 1k examples per class to create a manageable dataset for clustering, but did not specify how or which examples were selected. Their choice was made in the context of competing subspace clustering methods to which their model was being compared, many of which are not able to scale to larger datasets due to the memory constraints involved in computing a similarity matrix necessary for their methods. We tried without success to contact the authors to obtain the same subset of the data. Therefore, we decided to

Dataset	MNIST		FMNIST		SOP	
Metrics	ACC	NMI	ACC	NMI	ACC	NMI
K-means <sup>i</sup>	.572	.500	.474	.512	-	-
SC <sup>i</sup>	.696	.663	.508	.575	-	-
AC <sup>i</sup>	.695	.609	.500	.564	-	-
NMF <sup>i</sup>	.545	.608	.434	.425	-	-
DEC <sup>ii</sup>	.843	.772	.590	.601	<b>.229</b>	.121
DCN <sup>ii</sup>	.833	.809	.587	.594	.213	.084
JULE <sup>ii</sup>	.964	.913	.563	.608	-	-
VaDE <sup>ii</sup>	.945	.876	.578	.630	-	-
SpectralNet <sup>ii</sup>	<b>.971</b>	.924	.533	.552	-	-
DualAE <sup>ii</sup>	<b>.978</b>	<b>.941</b>	.662	.645	-	-
DAC <sup>ii</sup>	.966	<b>.967</b>	.615	.632	<b>.231</b>	.098
SSC <sup>iii</sup>	.430	.568	.359	.181	.127	.007
LLR <sup>iii</sup>	.552	.665	.345	.254	.224	<b>.173</b>
DSC-Net <sup>iii</sup>	.659	.730	.606	.617	<b>.269</b>	<b>.146</b>
KSSC <sup>iii</sup>	.585	.677	.382	.197	<b>.268</b>	<b>.152</b>
k-SCN <sup>iii</sup>	.871	.782	.638	.620	<b>.229</b>	<b>.166</b>
NCSC <sup>iii</sup>	.941	.861	<b>.721</b>	<b>.686</b>	<b>.275</b>	<b>.138</b>
ClustGAN <sup>iv</sup>	.950	.890	.630	.640	-	-
InfoGAN <sup>iv</sup>	.890	.860	.610	.590	.198	.082
DLS-Clst <sup>iv</sup>	<b>.975</b>	<b>.936</b>	<b>.693</b>	<b>.669</b>	-	-
IIC <sup>v</sup>	<b>.992</b>	<b>.978</b>	.657	.637	-	-
DCCM <sup>v</sup>	-	-	.657	.637	-	-
DCCS <sup>v</sup>	<b>.989</b>	<b>.970</b>	<b>.756</b>	<b>.704</b>	-	-
Ours	.943	.905	<b>.721</b>	<b>.691</b>	<b>.229</b>	.072
Ours (raw)	.877	.856	<b>.704</b>	<b>.690</b>	.204	.063

Table 1: Clustering performance results on 3 datasets w.r.t. ACC and NMI (top 5 in bold). ‘‘Ours’’ indicates HC-MGAN, with (raw) indicating no refinement operations performed. <sup>i</sup>: non-deep. <sup>ii</sup>: varied DC methods. <sup>iii</sup>: Subspace Clustering. <sup>iv</sup>: GAN-based DC. <sup>v</sup>: DC w/ data augmentation.

evaluate our results on the entire data. Due to the class imbalance, we compute both the mean ACC over classes (reported in the table) and the overall ACC achieved by HC-MGAN, which are, respectively, 0.229 and 0.221. NMI is invariant to class imbalance. On this dataset, our model does not perform as well as the subspace clustering group, especially NCSC, KSSC and DSC-Net, even though it was on par with NCSC and greatly surpassed the KSSC and DSC-Net on other datasets. But for other DC methods reported from (Zhang et al. 2019b), our method performs closely, even outperforming DCN and InfoGAN in accuracy.

**Qualitative Analysis** Figure 5 shows a visualization of the clustering tree for FMNIST. For each node  $k$  in the tree, it displays a grid of 25 real examples sampled from a multinational distribution over the dataset with weights for the  $i$ -th example given by  $s_{k,i}$ , yielding a representative view of  $s_k$ ’s soft cluster. Moreover, a color coded bar graph beside each grid displays the % of the global probability mass for each class  $A$  in  $s_k$ , given by  $100 \cdot \sum_i s_{k,i} \cdot \mathbf{1}(c_i = A) / \sum_i \mathbf{1}(c_i = A)$ , where  $\mathbf{1}$  is an indicator function and  $c_i$  is the class of the  $i$ -th example. Through a quick inspection, we observe that the 1<sup>st</sup> split ( $s_0$ ) occurred with almost perfect precision, with examples of the same class having their prob-

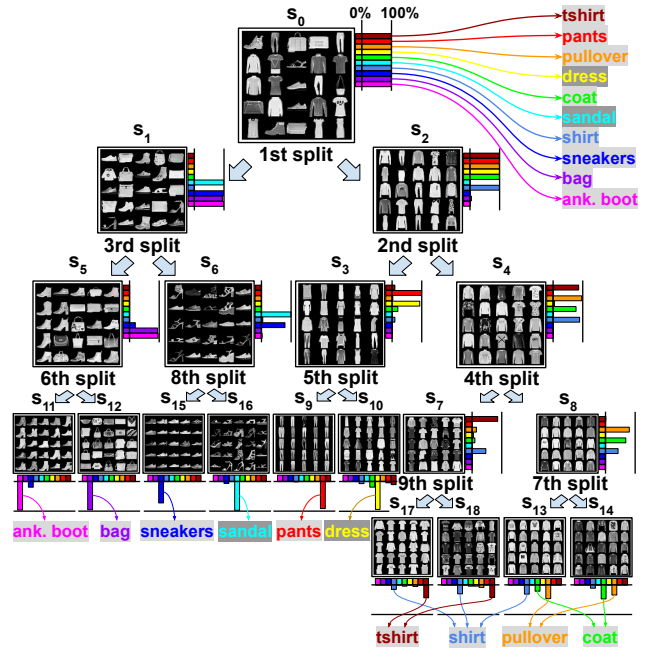


Figure 5: FMNIST cluster tree. Grids show sampled images; bar graphs beside  $k$ -th node’s grid represent % of total probability mass for each class in that node (best viewed in color).

ability mass nearly entirely allocated to either  $s_1$  or  $s_2$ . We begin to notice some imprecision at the 2<sup>nd</sup> split ( $s_2$ ), with a small portion of the probability mass of classes coat and t-shirt being assigned to  $s_3$  while the largest portion went to  $s_4$ , and at the 3<sup>rd</sup> split ( $s_1$ ), with a small portion of the sneakers’ mass being assigned to  $s_5$  and the largest portion going to  $s_6$ . The most imprecise splits occurred during the 4<sup>th</sup>, 7<sup>th</sup> and 9<sup>th</sup> splits (*i.e.*,  $s_4$ ,  $s_7$  and  $s_8$ ), but the classes involved in these splits (t-shirt, pullover, coat and shirt) are the most visually similar in the dataset, thus being the hardest to accurately separate into clusters. The other classes are relatively well separated. One fact that can explain why our method performed very well w.r.t. NMI on FMNIST is that most of the probability mass of each class for which our method exhibited low accuracy was assigned to at most two clusters. For the NMI metric, having the wrongfully assigned classes concentrated on fewer clusters provides better mutual information than having them spread throughout more clusters.

## Conclusion

We proposed a method for hierarchical clustering, leveraging the representation capacity of GANs/MGANs. To the best of our knowledge, this is the first hierarchical DC method. It creates a tree of clusters from top to bottom, where each leaf node defines a cluster. Clusters are divided in binary splits, performed in two phases: a raw split and a refinement. We have shown how well our method compares to other DC methods, which lack the advantages of hierarchical clustering, obtaining competitive benchmark results.

## Acknowledgments

The authors thank the partial support from Brazilian research supporting agencies: CNPq (grant PQ 313582/2018-1), FAPEMIG (grant CEX - PPM-00598-17, grant APQ-02337-21), and CAPES (grant 88887.506739/2020-00).

## References

- Achanta, R.; and Susstrunk, S. 2017. Superpixels and Polygons Using Simple Non-Iterative Clustering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Cai, D.; He, X.; Wang, X.; Bao, H.; and Han, J. 2009. Locality Preserving Nonnegative Matrix Factorization. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- Chang, J.; Wang, L.; Meng, G.; Xiang, S.; and Pan, C. 2017. Deep Adaptive Image Clustering. In *IEEE International Conference on Computer Vision (ICCV)*.
- Chen, X.; Duan, Y.; Houthoofd, R.; Schulman, J.; Sutskever, I.; and Abbeel, P. 2016. InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Ding, F.; and Luo, F. 2019. Clustering by Directly Disentangling Latent Space. arXiv:1911.05210.
- Elhamifar, E.; and Vidal, R. 2013. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*.
- Ghasedi Dizaji, K.; Herandi, A.; Deng, C.; Cai, W.; and Huang, H. 2017. Deep Clustering via Joint Convolutional Autoencoder Embedding and Relative Entropy Minimization. In *IEEE International Conference on Computer Vision (ICCV)*.
- Ghosh, A.; Kulharia, V.; Nambodiri, V.; Torr, P. H. S.; and Dokania, P. K. 2018. Multi-agent Diverse Generative Adversarial Networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Net. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Gowda, K. C.; and Krishna, G. 1978. Agglomerative clustering using the concept of mutual nearest neighbourhood. *Pattern Recognition*.
- Hoang, Q.; Nguyen, T.; Le, T.; and Phung, D. 2018. MGAN: training generative adversarial nets with multiple generators. In *International Conference on Learning Representations (ICLR)*.
- Ji, P.; Zhang, T.; Li, H.; Salzmann, M.; and Reid, I. 2017. Deep Subspace Clustering Networks. In *International Conference on Neural Information Processing Systems (NeurIPS)*.
- Ji, X.; Henriques, J. F.; and Vedaldi, A. 2019. Invariant Information Clustering for Unsupervised Image Classification and Segmentation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Jiang, Z.; Zheng, Y.; Tan, H.; Tang, B.; and Zhou, H. 2017. Variational Deep Embedding: An Unsupervised and Generative Approach to Clustering. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- Joulin, A.; Bach, F.; and Ponce, J. 2010. Discriminative clustering for image co-segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; and Aila, T. 2020. Training Generative Adversarial Networks with Limited Data. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Kundu, J. N.; Gor, M.; Agrawal, D.; and Babu, R. V. 2019. GAN-Tree: An Incrementally Learned Hierarchical Generative Framework for Multi-Modal Data Distributions. In *IEEE/CVF International Conference on Computer Vision (ICCV)*.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Liu, G.; Lin, Z.; Yan, S.; Sun, J.; Yu, Y.; and Ma, Y. 2013. Robust Recovery of Subspace Structures by Low-Rank Representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*.
- Liu, Z.; Lin, G.; Yang, S.; Feng, J.; Lin, W.; and Goh, W. L. 2018. Learning Markov Clustering Networks for Scene Text Detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- MacQueen, J. 1967. Some methods for classification and analysis of multivariate observations. In *Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*.
- Mukherjee, S.; Asnani, H.; Lin, E.; and Kannan, S. 2019. ClusterGAN: Latent Space Clustering in Generative Adversarial Networks. *AAAI Conference on Artificial Intelligence (AAAI)*.
- Oh Song, H.; Xiang, Y.; Jegelka, S.; and Savarese, S. 2016. Deep Metric Learning via Lifted Structured Feature Embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Patel, V. M.; and Vidal, R. 2014. Kernel sparse subspace clustering. In *IEEE International Conference on Image Processing (ICIP)*.
- Shaham, U.; Stanton, K. P.; Li, H.; Basri, R.; Nadler, B.; and Kluger, Y. 2018. SpectralNet: Spectral Clustering using Deep Neural Networks. In *International Conference on Learning Representations (ICLR)*.
- Shi, J.; and Malik, J. 2000. Normalized cuts and image segmentation. *Pattern Analysis and Machine Intelligence (PAMI)*.
- Wu, J.; Long, K.; Wang, F.; Qian, C.; Li, C.; Lin, Z.; and Zha, H. 2019. Deep Comprehensive Correlation Mining for Image Clustering. In *IEEE/CVF International Conference on Computer Vision (ICCV)*.



Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.

Xie, J.; Girshick, R.; and Farhadi, A. 2016. Unsupervised Deep Embedding for Clustering Analysis. In *International Conference on Machine Learning (ICML)*.

Yang, B.; Fu, X.; Sidiropoulos, N. D.; and Hong, M. 2017. Towards K-Means-Friendly Spaces: Simultaneous Deep Learning and Clustering. In *International Conference on Machine Learning (ICML)*.

Yang, J.; Parikh, D.; and Batra, D. 2016. Joint Unsupervised Learning of Deep Representations and Image Clusters. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yang, X.; Deng, C.; Zheng, F.; Yan, J.; and Liu, W. 2019. Deep Spectral Clustering Using Dual Autoencoder Network. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Zelnik-Manor, L.; and Perona, P. 2004. Self-Tuning Spectral Clustering. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Zhang, H.; Xu, S.; Jiao, J.; Xie, P.; Salakhutdinov, R.; and Xing, E. P. 2018. Stackelberg GAN: Towards Provable Minimax Equilibrium via Multi-Generator Architectures. *arXiv:1811.08010*.

Zhang, T.; Ji, P.; Harandi, M.; Hartley, R.; and Reid, I. 2019a. Scalable Deep k-Subspace Clustering. In *Asian Conference on Computer Vision (ACCV) 2018*.

Zhang, T.; Ji, P.; Harandi, M.; Huang, W.; and Li, H. 2019b. Neural Collaborative Subspace Clustering. In *International Conference on Machine Learning (ICML)*.

Zhao, J.; Lu, D.; Ma, K.; Zhang, Y.; and Zheng, Y. 2020. Deep Image Clustering with Category-Style Representation. In *European Conference on Computer Vision (ECCV)*.