



MOX-Report No. 33/2021

**Scalable balanced training of conditional generative  
adversarial neural networks on image data**

Lupo Pasini, M.; Gabbi, V.; Yin, J.; Perotto, S.; Laanait, N.

MOX, Dipartimento di Matematica  
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

[mox-dmat@polimi.it](mailto:mox-dmat@polimi.it)

<http://mox.polimi.it>

# Scalable balanced training of conditional generative adversarial neural networks on image data

Massimiliano Lupo Pasini<sup>1</sup>, Vittorio Gabbi<sup>2</sup>, Junqi Yin<sup>1</sup>,  
Simona Perotto<sup>3</sup>, Nouamane Laanait<sup>4</sup>

May 19, 2021

<sup>1</sup>Oak Ridge National Laboratory, Computational Sciences and Engineering Division  
Oak Ridge, TN 37831, USA

<sup>2</sup>Politecnico di Milano, Piazza L. da Vinci 32, I-20133 Milano, Italy

<sup>3</sup>MOX– Dipartimento di Matematica, Politecnico di Milano

Piazza L. da Vinci 32, I-20133 Milano, Italy

<sup>4</sup>Anthem Blue Cross, Atlanta, GA, 30326, USA

## Abstract

We propose a distributed approach to train deep convolutional generative adversarial neural network (DC-CGANs) models. Our method reduces the imbalance between generator and discriminator by partitioning the training data according to data labels, and enhances scalability by performing a parallel training where multiple generators are concurrently trained, each one of them focusing on a single data label. Performance is assessed in terms of inception score and image quality on MNIST, CIFAR10, CIFAR100, and ImageNet1k datasets, showing a significant improvement in comparison to state-of-the-art techniques to training DC-CGANs. Weak scaling is attained on all the four datasets using up to 1,000 processes and 2,000 NVIDIA V100 GPUs on the OLCF supercomputer Summit.

This manuscript has been authored in part by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

## 1 Introduction

Generative adversarial neural networks (GANs) [1] [2] [3] [4] are deep learning (DL) models whereby a dataset is used by an agent, called the generator, to sample white noise from a latent space and simulate a data distribution to create new (fake) data

that resemble the original data it has been trained on. Another agent, called the discriminator, has to correctly discern between the original data (provided by the external environment for training) and the fake data (produced by the generator). The generator prevails over the discriminator if the latter does not succeed in distinguishing anymore the original from the fake. The discriminator prevails over the generator if the fake data created by the generator is categorized as fake, and the original data is still categorized as original. An illustration that describes a GANs model is shown in Figure 1. Originally, GANs have been used on image data to improve the generalizability of DL models for object recognition. Image data are represented as pixels in a Cartesian structure. For each pixel, a set of values called channels are assigned to describe the local graphic properties. The channels per pixel are only one for black and white images, and colored images have three channels (red, green and blue). In general, the graphic variability between classes is more pronounced than the graphic variability within a specific class, because objects of the same type generally resemble more than objects of different nature. The benchmark datasets we consider are characterized by labels that clearly separate images according to their category, and the category is related to the type of object represented in the image.

The training of GANs is driven by the values of the cost functions associated with the agents [5]. The cost functions used to evaluate the performance of the discriminator and the generator are related to the number of false positives (original images identified by the discriminator as fake) and false negatives (fake images identified by the discriminator as original). The task of the discriminator is relatively simple in that it only has to assign a Boolean value to an image, according to whether the image is predicted as original or fake. On the contrary, the generator needs to map white noise sampled from the latent space into newly created images, and the images created must reproduce relevant features that belong to each data category represented in the training data.

This imbalance between the difficulty of the computational tasks of discriminator and generator is natural in GANs and cost functions currently used to measure the performance of the generator do not retain information about the disparity in computational tasks between discriminator and generator. As a result, the precision attained by the discriminator in performing its tasks (saying if an image is fake or original) is always higher than the precision with which the generator performs its own (create a whole set of fake images from white noise). Recent game theoretic results show that the unbalanced training of GANs can cause the generator to cycle [6] or converge to a (potentially bad) local optimum [7], which causes the generator to get stuck reproducing only one specific data point (this phenomenon is known in the DL literature as mode collapse). It is thus important to balance the training of GANs models in order to improve the performance of the generator and obtain fake images with similar features to the ones contained in the original data, but this task is challenging [8] [9] [10] [11] [12].

Some recent approaches have tackled the imbalance of the two agents by changing the numerical optimization used to train the GANs model. An example is the Competitive Gradient Descent method (CGD) [13], which recasts the GANs training as a zero-

sum game, whereby the discriminator and the generator compete against each other, and the goal is to identify an equilibrium between the agents. However, the zero-sum formulation does not reflect well the interaction between generator and discriminator during the GANs training, since the loss of one agent does not directly translate into the gain of the other agent, as it is indeed assumed in a zero-sum game. Other recent approaches have tackled the challenge of imbalance between discriminator and generator by improving the complexity of the GANs model [14] [15] [16] [17] [18] [19] [20] [21] [22]. Among these approaches, Conditional GANs (CGANs) [23] [24] [25] proceed by expanding the latent space used as input for the generator by adding information about the data categories. The role of CGAN models is to reconstruct a joint data distribution defined on the expanded latent space that combines the image data with the corresponding labels. The inclusion of data labels as an additional latent space variable facilitates the generator in discerning relevant features that make an image more likely to belong to a data category than to another. However, datasets with a large number of categories still pose a non-trivial challenge that prevents the generator from attaining a good performance in creating new fake images still due to a large variability between classes.

In addition, all the existing approaches to train GANs are characterized by a limited parallelizability, in that existing parallel techniques for GANs are based on data parallelization that distributes large-scale data to multiple replicas of the same model via ensemble learning, and do not further enhance the scalability of GANs by attempting any model parallelization [9]. Therefore, state-of-the-art GANs approaches do not fully leverage high-performance computing (HPC) facilities to attain a better performance.

*We propose a novel distributed approach to train CGANs through a non-zero-sum game formulation that uses data categories to address the performance disparity between discriminator and generator and improve the scalability of the CGANs training via model parallelization.* We use the labels to split the data and process each class independently, using a generator for each class. Our distributed approach relies on a factorization of the data distribution where each factor is associated with a single data category. The factorization of the probability distribution makes our approach differ from standard CGANs, where the joint data probability is never decomposed into simpler factors and a single generator is still assigned with the task of creating new images that span all the data categories. The data splitting performed according to the labels removes the variability between classes and thus corrects the imbalance of standard GANs training. Because of the independence of each generator from the others, the generators can be trained concurrently and this enhances scalability.

## **2 Distributed conditional generative adversarial neural networks**

Our approach aims at implementing a distributed version of CGANs as a non-zero-sum game. Figure 2 describes how our approach relates to the other generative models

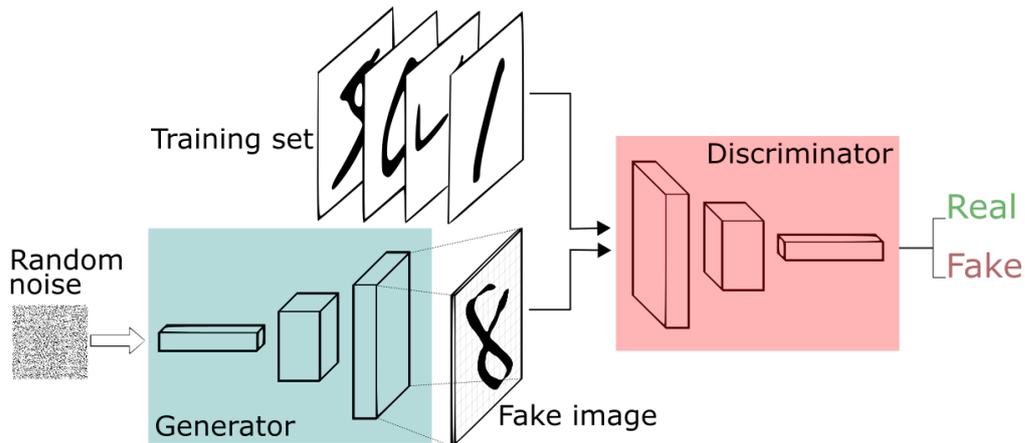


Figure 1: The GAN framework pits two adversaries against each other in a game. Each player is represented by a differentiable function controlled by a set of parameters. Typically these functions are implemented as deep neural networks. Training examples are randomly sampled from the training set and used as input for the first player, the discriminator. The goal of the discriminator is to output the probability that its input is real rather than fake, under the assumption that half of the inputs it is ever shown are real and half are fake. Image from <https://sthalles.github.io/intro-to-gans/>

presented in the literature. We first define the following input and output spaces, each with an associated probability distribution:

- $Z$  is a noise space used to seed the generative model.  $Z = \mathbb{R}^{d_Z}$ , where  $d_Z$  is a hyperparameter. Values  $\mathbf{z} \in Z$  are sampled from a noise distribution  $p_{\mathbf{z}}(\mathbf{z})$ . In our experiments  $p_{\mathbf{z}}$  is a white noise model.
- $Y$  is an embedding space used to condition the generative model on additional external information, drawn from the training data.  $Y = \mathbb{R}^{d_Y}$ , where  $d_Y$  is a hyperparameter. Using conditional information provided in the training data, we define a density model  $p_{\mathbf{y}}(\mathbf{y})$ .
- $X$  is the data space which represents an image output from the generator or input to the discriminator. Values are normalized pixel values:  $X = [0, 1]^W \times C$ , where  $W$  represents the resolution of the input images, and  $C$  is the set of distinct color channels in the input images. Using the images in the training data and their associated conditional data, we can define a density model  $p_{\text{data}}(\mathbf{x})$  of face images. This is exactly the density model we wish to replicate with the overall model in this paper.

We now define two functions:

- $G : Z \times Y \rightarrow X$  is the conditional generative model (or generator), which accepts noise data  $\mathbf{z} \in Z$  and produces an image  $\mathbf{x} \in X$  conditional to the external information  $\mathbf{y} \in Y$ .

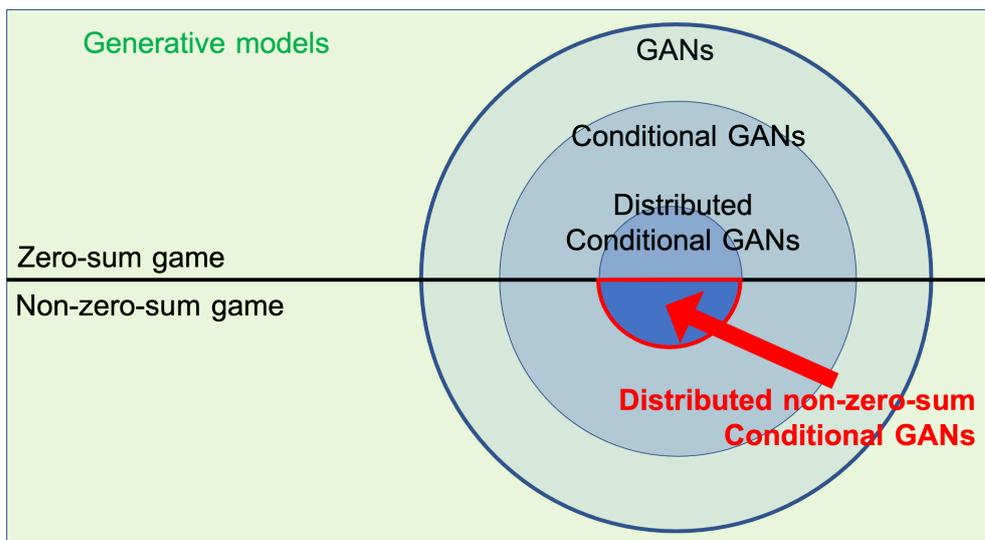


Figure 2: Scheme representing the relationship between the generative models currently developed in the deep learning community.

- $D : X \rightarrow [0, 1]$  is the discriminative model (or discriminator), which accepts an image  $\mathbf{x}$  and condition  $\mathbf{y}$  and predicts the probability under condition  $\mathbf{y}$  that  $\mathbf{x}$  came from the empirical data distribution rather than from the generative model.

The goal of CGANs is to provide a model that estimates the probability distribution  $p_{\text{model}}(\mathbf{x}, \boldsymbol{\theta}, \mathbf{y})$ , parameterized by parameters  $\boldsymbol{\theta}$  that describes the DL model. We then refer to the likelihood as the probability that the model assigns to the training data:  $\prod_{i=1}^m p_{\text{model}}(\mathbf{x}_i, \boldsymbol{\theta}, \mathbf{y})$ , for a dataset containing  $m$  training samples  $\mathbf{x}_i$ . Among the different types of generative models, GANs is a type of model that works via the principle of *maximum likelihood*. The principle of maximum likelihood aims at choosing the parameters  $\boldsymbol{\theta}$  for the DL model that maximize the likelihood of the training data

$$\boldsymbol{\theta}^* = \underset{\boldsymbol{\theta}}{\operatorname{argmax}}(p_{\text{model}}(\mathbf{x}, \boldsymbol{\theta}, \mathbf{y})) \quad (1)$$

Using condition information provided in the training data, we define a density model  $p_{\mathbf{y}}(\mathbf{y})$ . CGANs use the Bayes theorem to combine the conditional probability  $p_{\text{model}}(\mathbf{x}|\mathbf{y})$  and the density model  $p_{\mathbf{y}}(\mathbf{y})$  to yield the joint model probability  $p_{\text{model}}(\mathbf{x}, \mathbf{y})$ :

$$p_{\text{model}}(\mathbf{x}, \boldsymbol{\theta}, \mathbf{y}) = p_{\text{model}}(\mathbf{x}, \boldsymbol{\theta}|\mathbf{y})p_{\mathbf{y}}(\mathbf{y}). \quad (2)$$

While (2) cannot be expressed in closed analytical form, CGANs can be trained without needing to explicitly define a density function, because this type of generative model offers a way to train the model while interacting only indirectly with  $p_{\text{model}}(\mathbf{x}, \boldsymbol{\theta}, \mathbf{y})$ , usually by sampling from it.

The two players in the game are represented by two functions. The discriminator is defined by a function  $D$  that takes  $\mathbf{x}$  as input and uses  $\boldsymbol{\theta}^{(D)}$  as parameters. The

generator is defined by a function  $G$  that takes  $\mathbf{z}$  as input and uses  $\boldsymbol{\theta}^{(G)}$  as parameters. Both players have cost functions that are defined in terms of both players' parameters. A consensus has been reached in the literature about which cost functions fully describe the performance of the discriminator due to the simplicity of the discriminator's task [1] [26], whereas the complexity of the computational task of the generator still keeps different options open as to the cost function that better describes the generator's performance.

The cost used for the discriminator is

$$J^{(D)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)}) = -\frac{1}{2}\mathbb{E}_{\mathbf{x}\sim p_{\text{data}}}\log D(\mathbf{x}) - \frac{1}{2}\mathbb{E}_{\mathbf{z}}\log(1 - D(G(\mathbf{z}))) \quad (3)$$

The cost function we choose for the generator is

$$J^{(G)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)}) = \frac{1}{2}\mathbb{E}_{\mathbf{z}}\log D(G(\mathbf{z})). \quad (4)$$

This cost function  $J^{(G)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)})$  quantitatively describes the ability of the generator in tricking the discriminator so that the discriminator confuses fake images as real. This choice of  $J^{(G)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)})$  better reflects the goal of the generator as an individual agent, with respect to other alternatives that force a strong dependence of  $J^{(G)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)})$  on  $J^{(D)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)})$ , such as, for instance, in the zero-sum game formulation where  $J^{(G)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)}) = -J^{(D)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)})$ . The definition in (4) ignores the false positive, because the original images are not a product of the generator (only fake images are).

The discriminator wishes to minimize  $J^{(D)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)})$  and must do so while controlling only  $\boldsymbol{\theta}^{(D)}$ . The generator wishes to maximize  $J^{(G)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)})$  and must do so while controlling only  $\boldsymbol{\theta}^{(G)}$ . The solution of this mini-max game is a Nash equilibrium. Here, we use the terminology of local differential Nash equilibria [27]. In this context, a Nash equilibrium is a tuple  $(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)})$  that is a local minimum of  $J^{(D)}$  with respect to  $\boldsymbol{\theta}^{(D)}$  and a local maximum of  $J^{(G)}$  with respect to  $\boldsymbol{\theta}^{(G)}$ .

The training process consists of a numerical optimization scheme that iteratively updates  $\boldsymbol{\theta}^{(D)}$  and  $\boldsymbol{\theta}^{(G)}$  [9]. On each step, two minibatches are sampled: a minibatch of  $\mathbf{x}$  values from the dataset and a minibatch of  $\mathbf{z}$  values drawn from the model's prior over latent variables. The standard choice of a numerical optimization algorithm used to update  $\boldsymbol{\theta}^{(D)}$  and  $\boldsymbol{\theta}^{(G)}$  is a gradient-based optimization algorithm called Adam [28].

Our distributed approach to train CGANs relies on the the equality

$$p_{\text{model}}(\mathbf{x}, \boldsymbol{\theta}) = \sum_{k=1}^K p_{\text{model}}(\mathbf{x}, \boldsymbol{\theta}|\mathbf{y}_k)p_{\mathbf{y}}(\mathbf{y}_k) \quad (5)$$

to distribute the computation of each term  $p_{\text{model}}(\mathbf{x}, \mathbf{y}_k)$  by training  $K$  distributed CGANs, each one per class, and then we combine the results at the end of each training to yield  $p_{\text{model}}(\mathbf{x})$ . The numerical examples presented in this paper are characterized by a one-to-one mapping between  $\mathbf{y}_k$  and the labels in the image dataset. The advantage of our

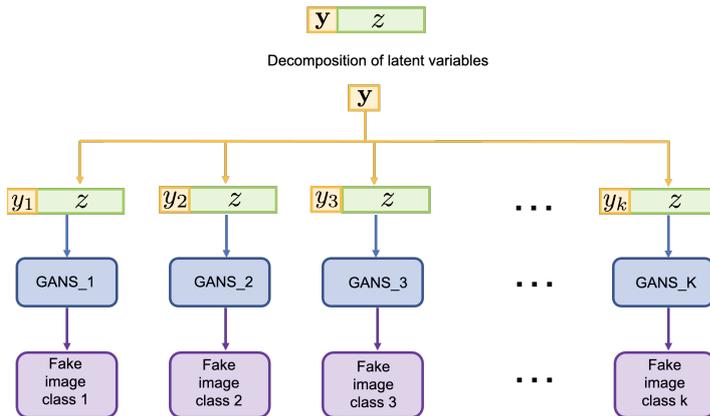


Figure 3: Illustration of distributed CGANs.

approach consists in the fact that all the  $K$  distributed CGANs can be trained concurrently and independently of each other, thus exposing the model to a higher level of parallelism. If the complexity representation of the objects in each category is comparable, the training time for each distributed CGANs model is approximately the same, which in turn translates into promising performance in terms of weak scalability (the time-to-solution is constant for an increasing number of processors used to solve problems of increasing size so that the computational workload per processor is unchanged). An illustration that describe the distribution of CGANs is provided in Figure 3.

### 3 Numerical results

In this section we compare the performance of standard deep convolutional GANs [2] (DC-GANs) with deep convolutional conditional GANs (DC-CGANs) [29] and our distributed approach to train DC-CGANs. The comparison is performed on a quantitative level by measuring the Inception Score (IS) [30]. The IS takes a list of images and returns a single floating point number, the score. The score is a measure of how realistic a GAN’s output is. IS is an automatic alternative to having humans grade the quality of images. The score measures two things simultaneously: the image variety (e.g. each image is a different breed of dog), and whether each image distinctly looks like a real object. If both things are true, the score will be high. If either or both are false, the score will be low. A higher score describes better performance for GANs, as it means that the GAN model can generate many different distinct and realistic images. The lowest score possible is zero. Mathematically the highest possible score is infinity, although in practice there will probably emerge a finite ceiling.

The specifics of the architectures for generator and discriminator used to build DC-GANs models are described in Tables 1 and 2. The specifics of the architectures for generator and discriminator used to build DC-CGANs models are described in Tables 3 and 4. The architecture of generator and discriminator for distributed DC-CGANs are

the same as for DC-GANs. The training is performed using the optimizer Adam and a learning rate of  $2e - 4$ , and a total number of 1,000 epochs for all the types of GANs we consider on each dataset.

### 3.1 Hardware description

The numerical experiments are performed using Summit [31], a supercomputer at the Oak Ridge Leadership Computing Facility (OLCF) at Oak Ridge National Laboratory. Summit has a hybrid architecture, and each node contains two IBM POWER9 CPUs and six NVIDIA Volta GPUs all connected together with NVIDIA’s high-speed NVLink. Each node has over half a terabyte of coherent memory (high bandwidth memory + DDR4) addressable by all CPUs and GPUs plus 1.6 TB of non-volatile memory (NVMe) storage that can be used as a burst buffer or as extended memory. To provide a high rate of communication and I/O throughput, the nodes are connected in a non-blocking fat-tree using a dual-rail Mellanox EDR InfiniBand interconnect.

### 3.2 Software description

The numerical experiments are performed using `Python3.7` with `PyTorch v1.3.1` package [32] for autodifferentiation to train the DL models with the use of GPUs, and the `mpi4py v3.0.2` tool is used for distributed computing.

As for the DC-GANs and the DC-CGANs approach, generator and discriminator are mapped to the same MPI processes. As for the distributed DC-CGANs, there are multiple discriminator-generator pairs, each one associated with a specific data class, and every discriminator-generator pair is mapped to an MPI process. Each MPI process instantiated in the distributed DC-CGANs is linked to two GPUs, one dedicated to training the discriminator and one dedicated to training the generator. Therefore, the total number of GPUs used with distributed DC-CGANs amounts to twice the number of MPI processes instantiated.

### 3.3 MNIST

The ANOVA test run on the MNIST dataset produces a value for the  $F$ -statistic equal to 38.78, which leads to a p-value close to zero. This indicates that there is strong statistical evidence to claim that the variability between images of different classes is larger than the variability between images of the same class. The IS obtained with all the three GANs models is shown in Table 5. The results show a similar performance for DC-GANs, DC-CGANs and distributed DC-CGANs, with a slight improvement using the latter over the other two. Samples of fake images generated by the generator of DC-GANs, DC-CGANs and distributed DC-CGANs trained on MNIST are shown in Figures 6, 7 and 8 respectively. The performance of DC-CGANs is very similar to the one of DC-GANs, because some of the digits are still written in a wobbly manner, preventing a clear understanding of what is the actual digit represented. This happens in situations where multiple digits resemble very much (the shape of a 3 is very similar

Generator						
Layer	Input dim	Output dim	Activation	Kernel size	stride	padding
Input	100	8192	ReLU	/	/	/
Resizing						
Batch normalization( epsilon = 1e-5 , momentum = 0.1)						
Upsample(scalefactor = 2)						
Convolution1	128	128	/	3	1	1
Batch normalization(epsilon = 0.8 , momentum = 0.1)						
leakyReLU(slope = 0.2, inplace=True)						
Upsample(scalefactor = 2)						
Convolution2	128	64	/	3	1	1
Batch normalization(epsilon = 0.8 , momentum = 0.1)						
leakyReLU(slope = 0.2, inplace=True)						
Convolution3	64	1 or 3	/	3	1	1
Tanh						

Table 1: Architecture of the generator in DC-GANs

Discriminator						
Layer	Input dim	Output dim	Activation	Kernel size	stride	padding
Convolution1	3	16	/	3	2	1
leakyReLU(slope = 0.2, inplace=True)						
Dropout(0.25)						
Convolution2	16	32	/	3	2	1
leakyReLU(slope = 0.2, inplace=True)						
Dropout(0.25)						
Batch normalization(epsilon = 0.8 , momentum = 0.1)						
Convolution3	32	64	/	3	2	1
leakyReLU(slope = 0.2, inplace=True)						
Dropout(0.25)						
Batch normalization(epsilon = 0.8 , momentum = 0.1)						
Convolution4	64	128	/	3	2	1
leakyReLU(slope = 0.2, inplace=True)						
Dropout(0.25)						
Batch normalization(epsilon = 0.8 , momentum = 0.1)						
output	2048	1	Sigmoid	/	/	/

Table 2: Architecture of the discriminator in DC-GANs

Generator						
Layer	Input dim	Output dim	Activation	Kernel size	stride	padding
Input	110	8192	ReLU	/	/	/
Resizing						
Batch normalization( epsilon = 1e-5 , momentum = 0.1)						
Upsample(scalefactor = 2)						
Convolution1	128	128	/	3	1	1
Batch normalization(epsilon = 0.8 , momentum = 0.1)						
leakyReLU(slope = 0.2, inplace=True)						
Upsample(scalefactor = 2)						
Convolution2	128	64	/	3	1	1
Batch normalization(epsilon = 0.8 , momentum = 0.1)						
leakyReLU(slope = 0.2, inplace=True)						
Convolution3	64	1 or 3	/	3	1	1
Tanh						

Table 3: Architecture of the generator in DC-CGANs

Discriminator						
Layer	Input dim	Output dim	Activation	Kernel size	stride	padding
Convolution1	3	16	/	3	2	1
leakyReLU(slope = 0.2, inplace=True)						
Dropout(0.25)						
Convolution2	16	32	/	3	2	1
leakyReLU(slope = 0.2, inplace=True)						
Dropout(0.25)						
Batch normalization(epsilon = 0.8 , momentum = 0.1)						
Convolution3	32	64	/	3	2	1
leakyReLU(slope = 0.2, inplace=True)						
Dropout(0.25)						
Batch normalization(epsilon = 0.8 , momentum = 0.1)						
Convolution4	64	128	/	3	2	1
leakyReLU(slope = 0.2, inplace=True)						
Dropout(0.25)						
Batch normalization(epsilon = 0.8 , momentum = 0.1)						
output	2048	1	None	/	/	/

Table 4: Architecture of the discriminator in DC-CGANs

	mean	standard deviation
DC-GANs	2.54	0.54
DC-CGANs	2.62	0.03
Distributed DC-CGANs	2.71	0.04

Table 5: Inception score associated with fake images generated by DC-GANs, DC-CGANs and distributed DC-CGANs trained on MNIST dataset. The inception score is calculated using 10 splits to collect statistics.

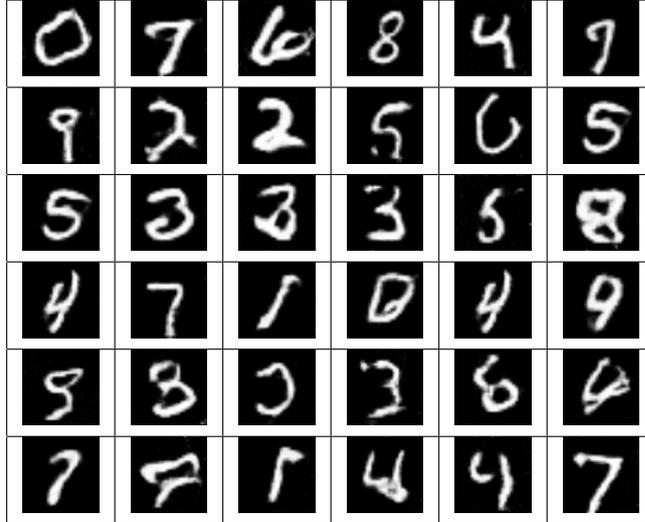


Table 6: Fake images generated by DC-GANs trained on MNIST dataset.

to the shape of an 8). The distributed DC-CGANs instead has a more consistent and clearer representation of all the digits.

### 3.4 CIFAR10

The ANOVA test run on the CIFAR10 dataset produces a value for the  $F$ -statistic equal to 22.12 for the red channel, 103.20 for the green channel, and 400.47 for the blue channel, which lead to p-values close to zero. Also in this case, there is strong statistical evidence to claim that the variability between images of different classes is larger than the variability between images of the same class. The IS obtained with all the three GANs models is shown in Table 9. The distributed approach for DC-CGANs improves the performance of the generator, as it is also shown by a comparison between samples of fake images generated by DC-GANs, DC-CGANs and distributed DC-CGANs models in Figures 10, 11 and 12 respectively. It can be noticed that the images produced with distributed DC-CGANs are more realistic than the ones produced by DC-GANs and standard DC-CGANs, as they represent objects that are easier to recognize as real.

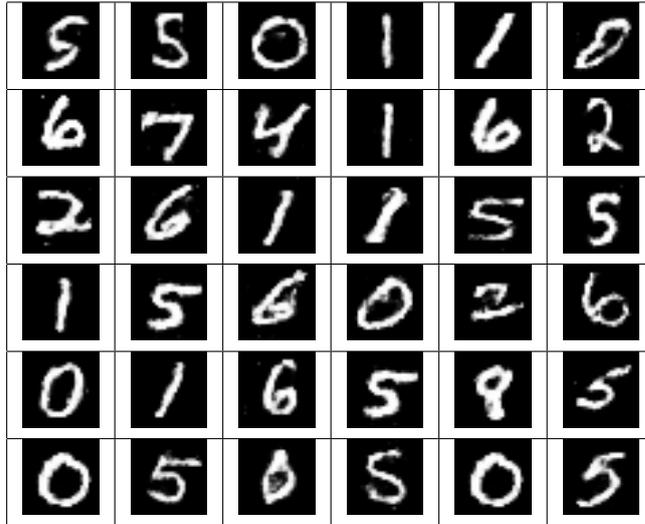


Table 7: Fake images generated by DC-CGANs trained on MNIST dataset.

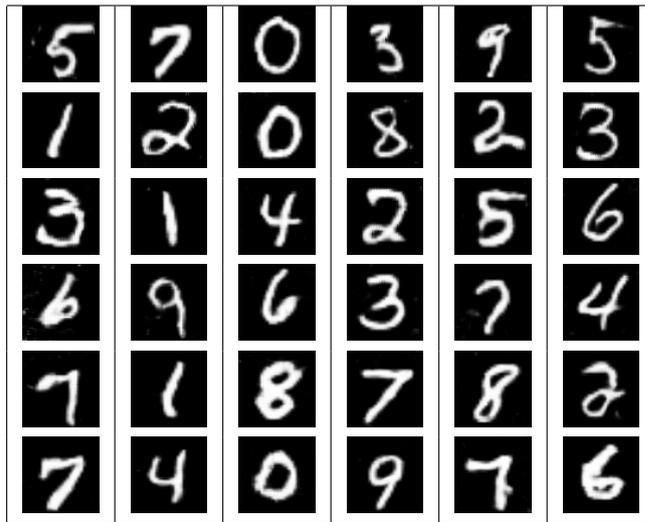


Table 8: Fake images generated by distributed DC-CGANs trained on MNIST dataset.

	mean	standard deviation
DC-GANs	4.39	0.28
DC-CGANs	5.69	0.31
Distributed DC-CGANs	6.43	0.25

Table 9: Inception score associated with images generated by DC-GANs, DC-CGANs and distributed DC-CGANs trained on CIFAR10 dataset. The inception score is calculated using 10 splits to collect statistics.

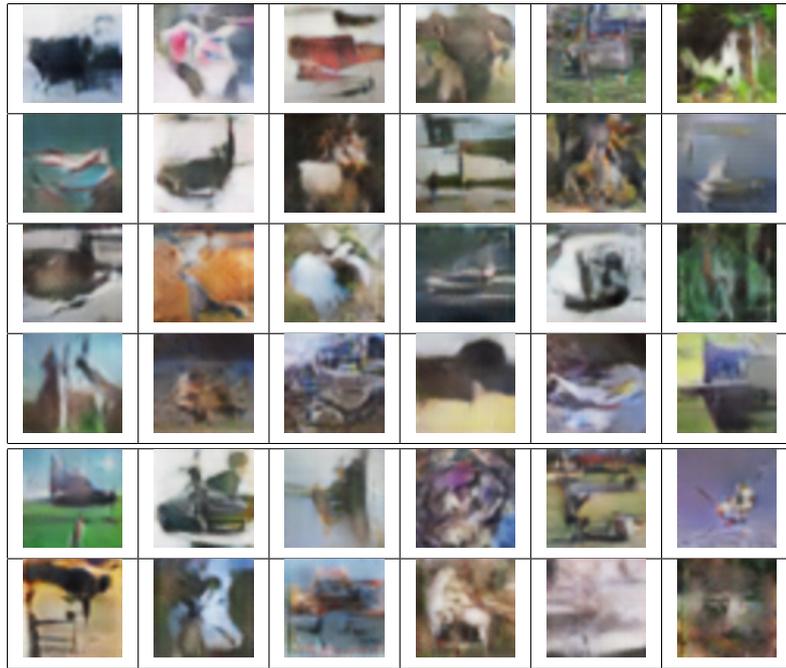


Table 10: Fake images generated by DC-GANs trained on CIFAR10 dataset.

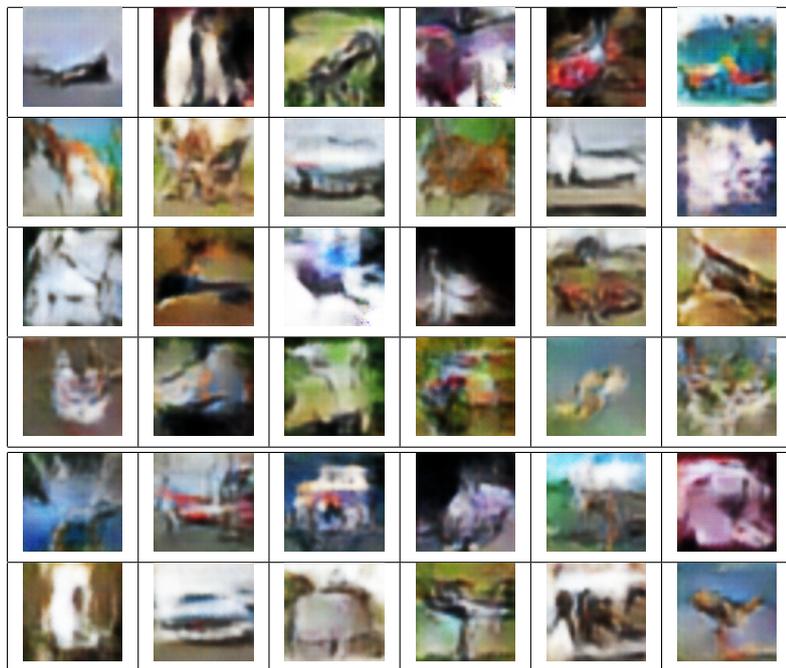


Table 11: Fake images generated by DC-CGANs trained on CIFAR10 dataset.

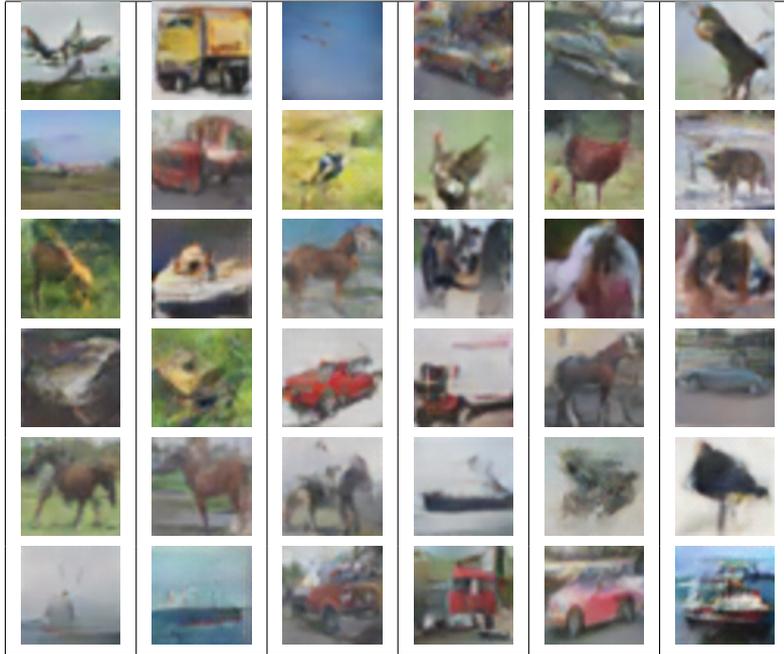


Table 12: Fake images generated by distributed DC-CGANs trained on CIFAR100 dataset.

### 3.5 CIFAR100

The ANOVA test run on the CIFAR100 produces a value for the  $F$ -statistic equal to 37.43 for the red channel, 23.74 for the green channel, and 26.08 for the blue channel, which lead to p-values close to zero. Also in this case, the variability between images of different classes is larger than the variability between images of the same class. The IS obtained with all the three GANs models is shown in Table 13. The results show an improved performance using DC-CGANs instead of DC-GANs, which is further improved by our distributed approach to train DC-CGANs. Samples of fake images generated by the generator of DC-GANs, DC-CGANs and distributed DC-CGANs trained on CIFAR100 are shown in Figures 14, 15 and 16 respectively, and the results still show distributed DC-CGANs generating more realistic pictures that can better be recognized as real objects.

### 3.6 ImageNet1k

The ANOVA test run on the ImageNet1k produces a value for the  $F$ -statistic equal to 57.43 for the red channel, 43.74 for the green channel, and 56.08 for the blue channel, which lead to p-values close to zero. The variability between images of different classes is larger than the variability between images of the same class, and this justifies a distributed training for GANs as we propose. The IS obtained with all the three GANs models is shown in Table 17. The results show an improved performance using DC-

	mean	standard deviation
DC-GANs	4.59	0.28
DC-CGANs	5.62	0.23
Distributed DC-CGANs	6.61	0.21

Table 13: Inception score associated with fake images generated by DC-GANs, DC-CGANs and distributed DC-CGANs trained on CIFAR100 dataset. The inception score is calculated using 10 splits to collect statistics.

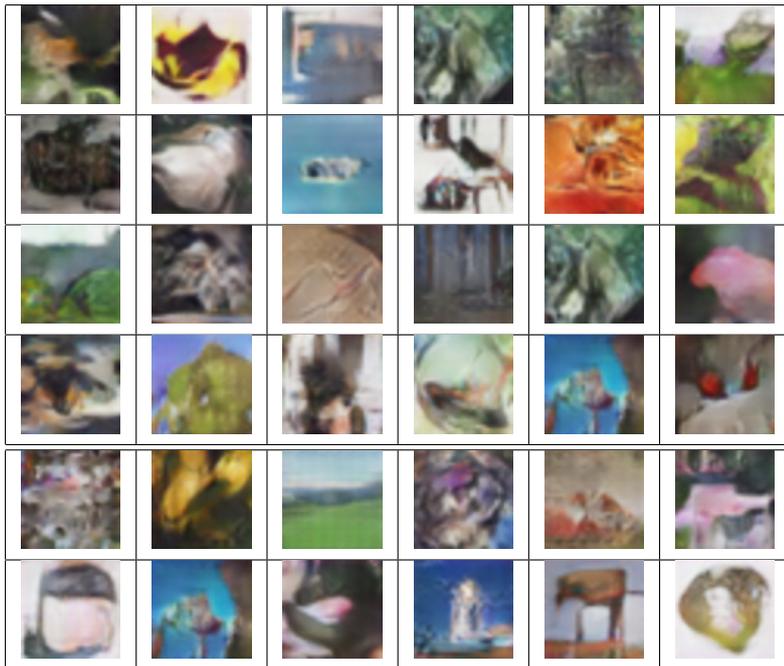


Table 14: Fake images generated by DC-GANs trained on CIFAR100.

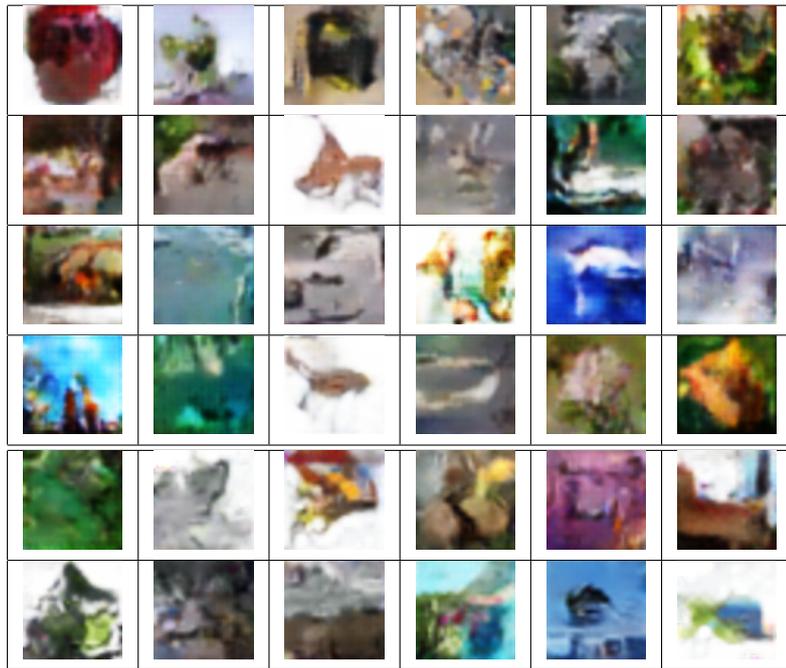


Table 15: Fake images generated by DC-CGANs trained on CIFAR100

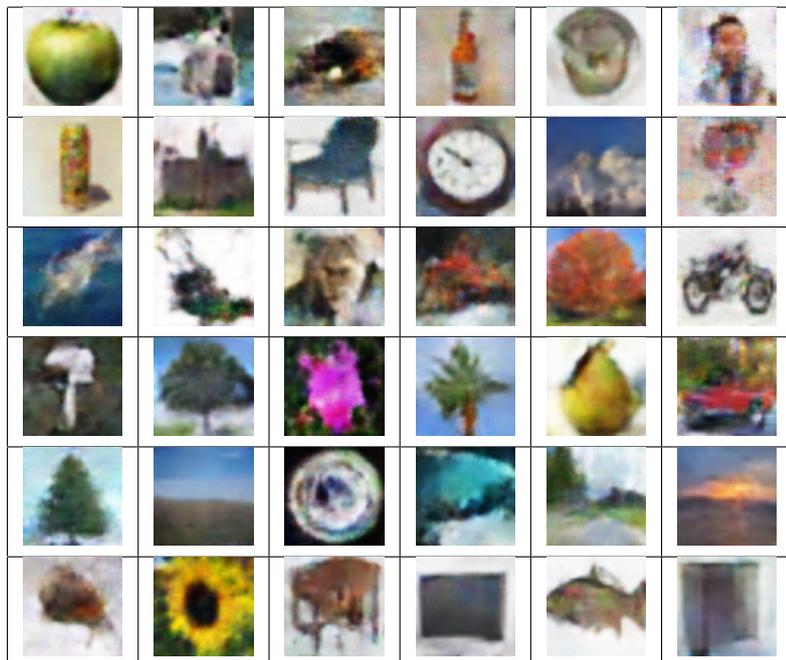


Table 16: Fake images generated by distributed DC-CGANs trained on CIFAR100.



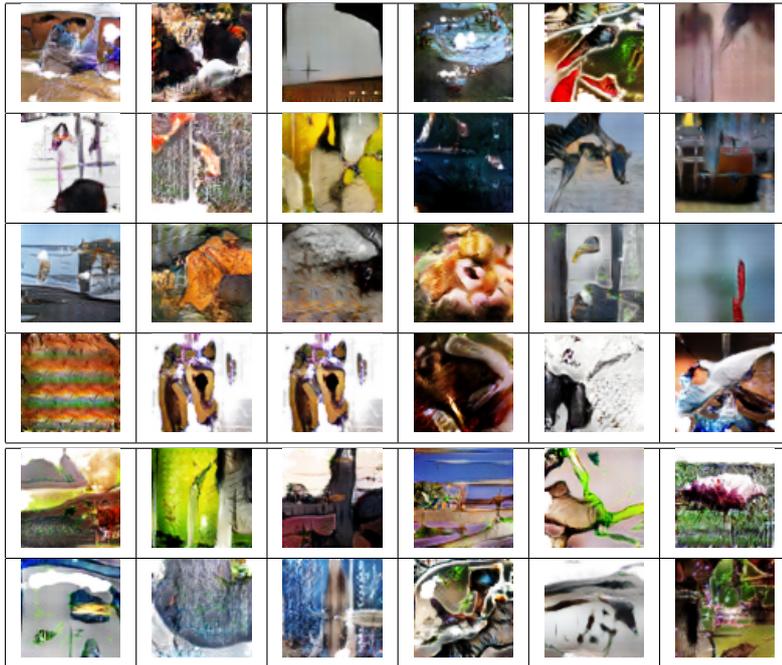


Table 19: Fake images generated by DC-CGANs trained on ImageNet1k

and the data classes. Each MPI process was mapped to two NVIDIA V100 GPUs, so that the neural networks for discriminator and generator for each data class would be trained on separate GPUs. For the MNIST and CIFAR10 datasets, the number of MPI processes spans the range from 1 to 10, whereas the number of MPI processes, for the CIFAR100 dataset was set to 10, 20, 40, 80 and 100, and for the ImageNet1k dataset was set to 100, 200, 400, 800 and 1,000. The results for the scalability tests are shown in Figure 4. The trend of the wall-clock time shows that the computational time to complete the training is not affected by the increasing number of data classes, as long as the computational workload for each MPI process stays fixed, thus showing that weak scaling is obtained by the training of distributed DC-CGANs on all three datasets.

The average GPU usage computed across all the GPUs engaged in the computation is 70% and the GPU memory occupation is around 31% for every numerical test performed.

## 4 Conclusions and future developments

We presented a distributed approach to train DC-CGANs models which reduces the imbalance between the computational tasks of the discriminator and the generator. The distribution of the task is based on partitioning the data in classes and it is justified by the variability of features between data classes that prevails over the variability of features within data points of the same class. The reduced imbalance results in better

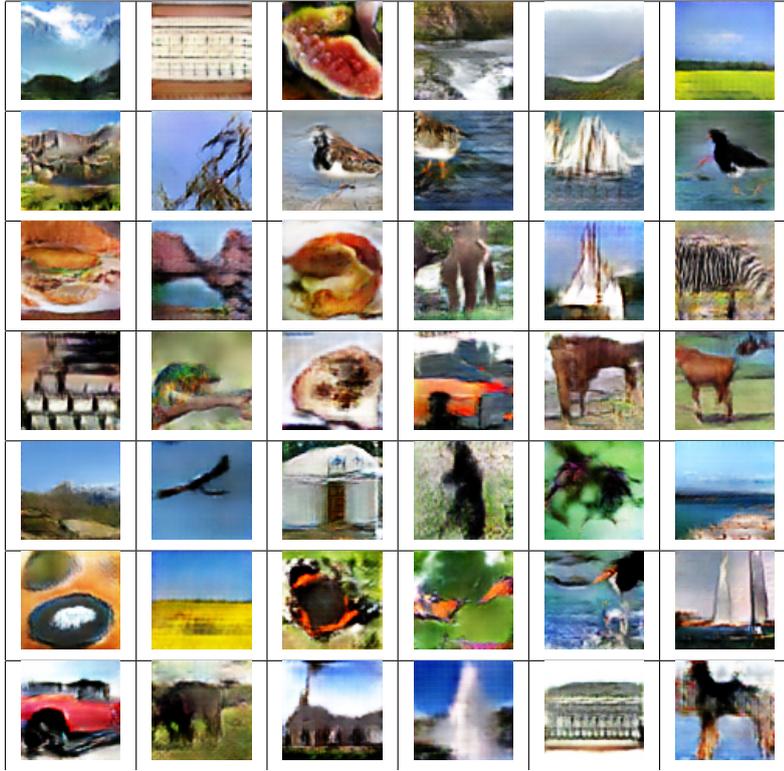


Table 20: Fake images generated by distributed DC-CGANs trained on ImageNet1k.

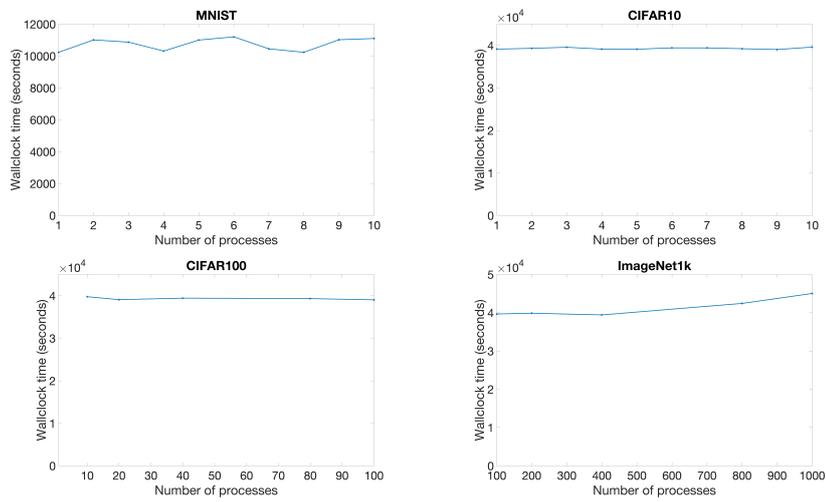


Figure 4: Results of weak scaling tests for MNIST, CIFAR10, CIFAR100, and ImageNet1k

images generated by the trained DC-CGANs compared to similar state-of-the-art approaches. The results have also shown an almost ideal weak scaling, indicating that HPC computing resources are fully leveraged.

Our future work aims at extending the study performed in this work to more complex but also more accurate neural network architectures, such as auxiliary classifier GANs (AC-GANs) [17], residual neural networks (ResNet) [18], self-attention generative adversarial neural networks (SAGANs) [20], and Wasserstein generative adversarial networks (WGANs) [33].

## Acknowledgements

Massimiliano Lupo Pasini thanks Dr. Vladimir Protopopescu for his valuable feedback in the preparation of this manuscript.

Research completed through the Artificial Intelligence Initiative sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U. S. Department of Energy.

This work used resources of the Oak Ridge Leadership Computing Facility, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

## Conflict of interest

The authors declare that they have no conflict of interest.

## References

- [1] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks. *arXiv:1406.2661 [cs, stat]*, June 2014. arXiv: 1406.2661.
- [2] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv:1511.06434 [cs]*, January 2016. arXiv: 1511.06434.
- [3] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved Techniques for Training GANs. *arXiv:1606.03498 [cs]*, June 2016. arXiv: 1606.03498.
- [4] Dimitri Bertsekas. Multiagent Rollout Algorithms and Reinforcement Learning. *arXiv:1910.00120 [cs]*, September 2019. arXiv: 1910.00120 version: 1.
- [5] Ian Goodfellow. NIPS 2016 Tutorial: Generative Adversarial Networks. *arXiv:1701.00160 [cs]*, April 2017. arXiv: 1701.00160.

- [6] Panayotis Mertikopoulos, Christos Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. *Proceedings of the 2018 Annual ACM-SIAM Symposium on Discrete Algorithms*, 2017.
- [7] Elad Hazan, Karan Singh, and Cyril Zhang. Learning linear dynamical systems via spectral filtering. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 6702–6712. Curran Associates, Inc., 2017.
- [8] Yann N. Dauphin, Harm de Vries, and Yoshua Bengio. Equilibrated adaptive learning rates for non-convex optimization. *arXiv:1502.04390 [cs]*, February 2015. arXiv: 1502.04390 version: 1.
- [9] Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv:1609.04747 [cs]*, June 2017. arXiv: 1609.04747.
- [10] John Duchi, Elad Hazan, and Yoram Singer. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. page 39.
- [11] Rachel Ward, Xiaoxia Wu, and Leon Bottou. AdaGrad stepsizes: Sharp Convergence over Nonconvex Landscapes, from any Initialization. *arXiv:1806.01811 [cs, stat]*, June 2018. arXiv: 1806.01811 version: 1.
- [12] Yang Zhao, Chunyuan Li, Ping Yu, Jianfeng Gao, and Changyou Chen. Feature Quantization Improves GAN Training. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11376–11386. PMLR, 13–18 Jul 2020.
- [13] Florian Schäfer and Anima Anandkumar. Competitive Gradient Descent. *arXiv:1905.12103 [cs, math]*, June 2020. arXiv: 1905.12103.
- [14] Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, and Y. Bengio. How to Construct Deep Recurrent Neural Networks. December 2013.
- [15] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going Deeper with Convolutions. *arXiv:1409.4842 [cs]*, September 2014. arXiv: 1409.4842.
- [16] Harikrishnan Sreekumaran, Ashish R. Hota, Andrew L. Liu, Nelson A. Uhan, and Shreyas Sundaram. Multi-Agent Decentralized Network Interdiction Games. *arXiv:1503.01100 [math]*, July 2015. arXiv: 1503.01100.
- [17] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional Image Synthesis with Auxiliary Classifier GANs. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*,

- volume 70 of *Proceedings of Machine Learning Research*, pages 2642–2651, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [18] M. Wang, H. Li, and Fang Li. Generative Adversarial Network based on Resnet for Conditional Image Restoration. *ArXiv*, abs/1707.04881, 2017.
- [19] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive Growing of GANs for Improved Quality, Stability, and Variation. *arXiv:1710.10196 [cs, stat]*, February 2018. arXiv: 1710.10196.
- [20] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention Generative Adversarial Networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 7354–7363. PMLR, 09–15 Jun 2019.
- [21] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large Scale GAN Training for High Fidelity Natural Image synthesis. In *International Conference on Learning Representations*, 2019.
- [22] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable Augmentation for Data-Efficient GAN Training, 2020. arXiv: 2006.10738.
- [23] Mehdi Mirza and Simon Osindero. Conditional Generative Adversarial Nets. *arXiv:1411.1784 [cs, stat]*, November 2014. arXiv: 1411.1784.
- [24] Takeru Miyato and Masanori Koyama. CGANs with Projection Discriminator. In *International Conference on Learning Representations*, 2018.
- [25] Ilya Kavalerov, Wojciech Czaja, and Rama Chellappa. cGANs with Multi-Hinge Loss. *CoRR*, abs/1912.04216, 2019.
- [26] Mario Lucic, Karol Kurach, Olivier Bousquet, and Sylvain Gelly. Are GANs Created Equal? A Large-Scale Study. *arXiv:1711.10337 [stat.ML]*, October 2018. arXiv: 1711.10337.
- [27] L. J. Ratliff, S. A. Burden, and S. S. Sastry. Characterization and Computation of Local Nash equilibria in Continuous Games. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 917–924, 2013.
- [28] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*, January 2017. arXiv: 1412.6980.
- [29] Jon Gauthier. Conditional generative adversarial nets for convolutional face generation - Technical report, Stanford University. 2015.
- [30] Shane Barratt and Rishi Sharma. A Note on the Inception Score. *arXiv:1801.01973 [cs, stat]*, June 2018. arXiv: 1801.01973.

- [31] Summit - Oak Ridge National Laboratory's 200 petaflop supercomputer <https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit/>.
- [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8026–8037. Curran Associates, Inc., 2019.
- [33] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein Generative Adversarial Networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.

## MOX Technical Reports, last issues

Dipartimento di Matematica  
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 31/2021** Ferraccioli, F.; Arnone, E.; Finos, L.; Ramsay, J.O.; Sangalli, L.M.  
*Nonparametric density estimation over complicated domains*
- 32/2021** Sangalli, L.M.  
*Spatial regression with partial differential equation regularization*
- 30/2021** Fumagalli, I.  
*A reduced 3D-0D FSI model of the aortic valve including leaflets curvature*
- 29/2021** Fumagalli, I.; Vitullo, P.; Scrofani, R.; Vergara, C.  
*Image-based computational hemodynamics analysis of systolic obstruction in hypertrophic cardiomyopathy*
- 28/2021** Ferro, N.; Perotto, S.; Bianchi, D.; Ferrante, R.; Mannisi, M.  
*Design of cellular materials for multiscale topology optimization: application to patient-specific orthopedic devices*
- 26/2021** Vigano, L.; Sollini, M.; Ieva, F.; Fiz, F.; Torzilli, G.  
*Chemotherapy-Associated Liver Injuries: Unmet Needs and New Insights for Surgical Oncologists*
- 27/2021** Scimone, R.; Menafoglio, A.; Sangalli, L.M.; Secchi, P.  
*A look at the spatio-temporal mortality patterns in Italy during the COVID-19 pandemic through the lens of mortality densities*
- 25/2021** Tenderini, R.; Pagani, S.; Quarteroni, A.; Deparis S.  
*PDE-aware deep learning for inverse problems in cardiac electrophysiology*
- 24/2021** Regazzoni, F.; Chapelle, D.; Moireau, P.  
*Combining Data Assimilation and Machine Learning to build data-driven models for unknown long time dynamics - Applications in cardiovascular modeling*
- 22/2021** Domanin, M.; Bennati, L.; Vergara, C.; Bissacco, D.; Malloggi, C.; Silani, V.; Parati, G.; Trimani, A.  
*Fluid structure interaction analysis to stratify the behavior of different atheromatous carotid plaques*