

Are conditional GANs explicitly conditional?

Houssemeddine Boulahbal^{1,2}
houssemeddine.boulahbal@renault.com

Adrian Voicila¹
adrian.voicila@renault.com

Andrew I. Comport²
Andrew.Comport@cnrs.fr

¹ Renault Software Factory,
2600 Route des Crêtes,
06560 Valbonne, France

² CNRS-I3S, Université Côte d'Azur,
2000 Route des Lucioles,
06903 Sophia Antipolis, France

Abstract

This paper proposes two important contributions for conditional Generative Adversarial Networks (cGANs) to improve the wide variety of applications that exploit this architecture. The first main contribution is an analysis of cGANs to show that they are not explicitly conditional. In particular, it will be shown that the discriminator and subsequently the cGAN does not automatically learn the conditionality between inputs. The second contribution is a new method, called *a contrario* cGAN, that explicitly models conditionality for both parts of the adversarial architecture via a novel *a contrario* loss that involves training the discriminator to learn unconditional (adverse) examples. This leads to a novel type of data augmentation approach for GANs (*a contrario* learning) which allows to restrict the search space of the generator to conditional outputs using adverse examples. Extensive experimentation is carried out to evaluate the conditionality of the discriminator by proposing a probability distribution analysis. Comparisons with the cGAN architecture for different applications show significant improvements in performance on well known datasets including, semantic image synthesis, image segmentation, monocular depth prediction and "single label"-to-image using different metrics including Fréchet Inception Distance (FID), mean Intersection over Union (mIoU), Root Mean Square Error log (RMSE log) and Number of statistically-Different Bins (NDB).

Since the seminal work in 2014, Generative Adversarial Networks (GANs) [14] have introduced an alternative framework for training generative models that has led to a multitude of high impact publications over a very large number of applications. Conditional GANs, introduced shortly after [35], have extended GANs to incorporate conditional information as input and have demonstrated resounding success for many computer vision tasks such as image synthesis [6, 22, 33, 41, 48, 49, 54], video synthesis [5, 32, 53], image correction[29, 42, 58], text-to-image[30, 43, 57, 60]. In all these works, the underlying GAN model as proposed in [14] and [35] have formed the basis for more advanced architectures and their properties have been analysed in detail and established in terms of convergence[27, 38], mode collapse[47], Nash equilibrium[12, 51], vanishing gradients[1], etc.

The « conditionality » of cGANs is at the crux of their theoretical contribution and its impact therefore merits in-depth analysis. From the existing literature it is not clear, however, if this now widely used architecture explicitly models or even learns conditionality. Empirically, the impressive results obtained with cGANs show that the generator automatically seeks to incorporate conditional variables into its generated output. Fundamentally the generator is, however, free to generate whichever output as long as it satisfies the discriminator. Therefore, the conditionality of cGAN also depends on the conditionality of the discriminator. This begs the question as to whether or not the baseline architecture of cGANs explicitly

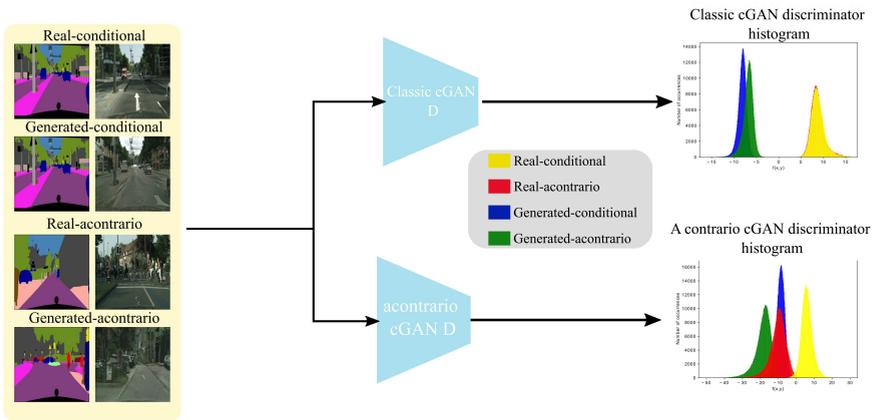


Figure 1: The classic cGAN and the proposed *a contrario* cGAN discriminators are tested with 500 validation images of the Cityscapes dataset on both conditional and unconditional label-to-image input set. Unconditional inputs (Real *a contrario* and Generated *a contrario*) set are obtained by randomly shuffling the original conditional sets of data. The classic cGAN discriminator fails to classify unconditional input set as false as seen by the histogram distributions on the right (real *a contrario* in red is classified as true). The proposed method trains the discriminator with a general *a contrario* loss to classify unconditional input set as fake (note that no extra training samples are required). The proposed *a contrario* cGAN correctly classifies all four modalities (blue, green, red, yellow) correctly.

models conditionality and if not, how can the core adversarial architecture be redefined to explicitly model conditionality? This is therefore the object of this paper.

Problems with cGANs conditionality have been observed independently for different tasks in the literature. Label-to-image tasks observe that using only adversarial supervision yields bad quality results [33, 41, 48, 54]. "Single Label"-to-image tasks [4] observe class leakage. It is also well known that cGANs are prone to mode collapse [44]. In this paper it is suggested that all these problems are related to the lack of a conditional discriminator.

Consider a simple test of conditionality on a learnt discriminator for the task of label-to-image translation, shown in Figure 1. The conditional label input is purposely swapped with a non-corresponding input drawn randomly from the input set (e.g. labels). From this test it is revealed that the discriminator does not succeed to detect the entire set of *a contrario* examples (defined in Section 2.3) as false input pairs. More extreme cases are shown in A.2. This suggests that the generator is not constrained by the discriminator to produce conditional output but rather to produce any output from the target domain (images in this case). Furthermore, in practice, the large majority of methods that exploit cGANs for label-to-image translation, if not all, add additional loss terms to the generator to improve conditionality. These loss terms are, however, not adversarial in their construction. For example, high resolution image synthesis approaches such as [54] suffer from poor image quality when trained with only adversarial supervision [48]. Considering the well known pix-to-pix architecture [22], a L1 loss was introduced to improve performance. This additional term seeks to enforce conditionality on the generator but does not act explicitly on the discriminator. Subsequently, one could question if the conditionality obtained by such methods is obtained via this loss term which is not part of the adversarial network architecture. Moreover, adding an extra loss term to the generator has now become the defacto method for improving cGANs results. For example perceptual loss [23] and feature-matching [45] have been proposed and reused by many others [6, 39, 41, 54]. As demonstrated in the experiments different tasks such as image-to-depth or image-to-label also exhibit these drawbacks.

In this paper it will be argued that simply providing condition variables as input is insuf-

ficient for modelling conditionality and that it is necessary to explicitly enforce dependence between variables in the discriminator. It will be demonstrated that the vanilla cGAN approach is not explicitly conditional via probabilistic testing of the discriminator’s capacity to model conditionality. With this insight, a new method for explicitly modelling conditionality in the discriminator and subsequently the generator will be proposed. This new method not only offers a solution for conditionality but also provides the basis for a general data augmentation method by learning from the contrary (*a contrario* data augmentation).

1 Related work

Conditional GANs Generative adversarial networks [14] can be considered as a zero-sum game between two player neural networks G and D competing to reach a Nash equilibrium. This game is commonly formulated through a min-max optimization problem as follows:

$$\min_{G \in \mathbb{G}} \max_{D \in \mathbb{D}} V(G, D) \quad (1)$$

where \mathbb{G} and \mathbb{D} are respectively the generator and discriminator function sets. A GAN is considered conditional [35] when the generator’s output is conditioned by an extra input variable. The condition variable can be any kind of information such as a segmentation mask, depth map, image or data from other modalities.

There are various methods that have been proposed for incorporating conditional information into the generator [10, 20, 41, 52, 66, 66]. Recently [49] introduced a classification-based feature learning module to learn more discriminating and class-specific features. Additional generator losses have also been proposed including feature matching [45], perceptual loss [23] and cycle-consistency loss [67]. While all these methods improve the conditionality of the generator, they assume that the discriminator models conditionality.

Alternatively several methods have been proposed to incorporate conditional information into the discriminator. [35] proposed an early fusion approach by concatenating the condition vector to the input of the discriminator. [24, 26, 33, 36, 39] proposed a late fusion by encoding the conditional information and introducing it into the final layers of the discriminator. [48] replaces the discriminator with a pixel-wise semantic segmentation network. Most of these methods are task specific and not general [24, 26, 48]. The method proposed in this paper is a task agnostic solution for incorporating any conditional variables and it will be shown to perform on several different datasets.

One concern raised in the literature is that improving either player is at the detriment to learning and can lead to mode collapse or vanishing gradient [1, 9]. In the present paper, whilst the aim is to define a conditional discriminator, it is argued that not only this does not affect the Nash equilibrium but also helps to avoid mode collapse and vanishing gradient. A simple example is that the proposed discriminator imposes conditionality and therefore the different condition inputs cannot be collapsed into a single mode.

Data augmentation for GANs Data augmentation is a highly useful strategy to improve generalization for machine learning that involves increasing the diversity of training data without collecting new samples. Augmentation strategies have involved performing a variety of operations on input training data including random cropping, horizontal flipping, various hand-crafted operations [13, 17] or even learning augmentation strategies from data [8, 62]. Performing data augmentation on GANs is less straightforward and has been limited to a handful of operations such as flipping, cropping and jittering [22]. [59, 64] propose a method to enforce the discriminator to remain unchanged by arbitrary augmentation. [25, 50, 63, 65] study the artifacts produced by augmentation and propose to perform augmentation on both real and generated images for both generator and discriminator training. Further, [25] proposes a method called stochastic discriminator augmentation that adjusts the augmentation

strength adaptively. Alternatively to all these approaches, the *a contrario* learning proposed in this paper performs data augmentation by providing false counter examples rather than augmenting the real data-set with transformations. The proposed data augmentation approach is therefore specific to adversarial architectures, however, it will be shown that by discriminating against undesirable distributions, the generator is better constrained towards the target distribution. In the specific case of conditionality, by providing *a contrario* examples to the discriminator as fakes, the discriminator constrains the generator to only learn conditional generations and the generator is therefore better constrained. It should be noted that in parallel to this work, the idea of augmenting the discriminator with fake examples was recently proposed in [46], however, this method did not consider discriminator conditionality and only investigated data augmentation through creating fake data.

Evaluation criteria for GANs Defining an evaluation criteria to assess and compare GAN architectures has been a major problem in advancing the field. The difficulty and advantage of GAN architectures lies in that the learnt discriminator replaces the classic objective function. Classic approaches for evaluating GANs often use the FID [18] or Inception scores [45] to evaluate high-level generation quality and diversity. For example, the inception score only allows to evaluate a perceptual image quality and does not account for consistency between generated output and label. By construction it is therefore domain specific and does not evaluate conditionality in the output. More recent approaches investigate more general metrics [2, 15, 19, 21, 61]. Nevertheless, none of these approaches specifically allow to evaluate the conditionality of conditional GANs. Such an evaluation metric should take both the condition variables and generated variables into account. Therefore, for the sake of evaluation, experiments were done on tasks that could provide a strong conditional metric to evaluate conditionality. The mIoU was used specifically for evaluating semantically conditional examples and log RMSE was considered for geometric depth conditionality.

2 Method

2.1 Classic cGAN

Classical cGAN training is based on conditionally paired sets of data $\mathcal{C}(\mathbf{x}, \mathbf{y})$ where $\mathbf{x} \sim p(\mathbf{x})$ is the condition variable and $\mathbf{y} \sim p(\mathbf{y}|\mathbf{x})$ is the real training variable. The generator of a cGAN outputs a transformed set of data $\mathcal{C}_G(\mathbf{x}, \mathbf{y}_G)$ composed of the generator output variable $\mathbf{y}_G \sim p_G(\mathbf{y})$ and the condition variable. These sets of data will be called "real-conditional" and "generated-conditional" respectively. The discriminator is defined as:

$$D(\mathbf{x}, \mathbf{y}) := \mathcal{A}(f(\mathbf{x}, \mathbf{y})) \quad (2)$$

Where $f(\cdot)$ is a neural network function of \mathbf{x} and \mathbf{y} , and \mathcal{A} is the activation function whose choice depends on the objective function. The cGAN objective function is defined as:

$$\mathcal{L}_{adv} = \min_G \max_D \left(\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}), \mathbf{y} \sim p(\mathbf{y}|\mathbf{x})} [\log(D(\mathbf{x}, \mathbf{y}))] + \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\log[1 - D(\mathbf{x}, G(\mathbf{x}))]] \right) \quad (3)$$

The min-max activation function is defined as a Sigmoid $\mathcal{A}(\mathbf{x}) = \left(\frac{1}{1+e^{-\mathbf{x}}} \right)$.

2.2 Evaluating conditionality

The objective of this section is to propose methods to test the conditionality of cGAN networks. As mentioned in Section 1, state-of-the-art approaches have focused on evaluating cGAN architectures with metrics applied to the generator output. Since the generator and discriminator are coupled, these metrics essentially evaluate the full GAN architecture.

A proposal is made to test the conditionality by visualizing the probability distribution at the output of the discriminator. Due to the fact that adversarial training involves a zero-sum game between a generator and a discriminator, both the generator and discriminator should seek to reach an equilibrium (Eq 1) at the end of training. One issue for GANs is that when the discriminator dominates there is a vanishing gradient problem [1]. It is therefore more difficult (but not impossible) to isolate the discriminator during training to evaluate its capacity to detect unconditional examples as false. For this reason, an optimal discriminator can be used to give insight for evaluation purposes as in [1, 12]. An optimal discriminator is essentially a binary classifier which classifies between true and fake (see Eq (3)).

In order to test the optimal discriminator, consider that the generator has been fixed after a certain number of iterations (the generator is initially random) and the discriminator has been allowed to converge to an optimal solution based on the following objective function:

$$\max_{D \in \mathbb{D}} V(G_{fixed}, D) \quad (4)$$

The evaluation subsequently involves analysing the distributions produced by the optimal discriminator (Eq (4)) given test distributions containing unconditional or *a contrario* sets of data-pairings. The capacity of the discriminator to correctly classify unconditional data as false is then analysed statistically. Section 2.3 provides a formal definition of these unconditional data pairings. Probability distributions are visualised and evaluated by histogram analysis on the discriminator features in the last convolution layer. Refer to Section A.1 for a more theoretical analysis and detailed results.

2.3 A contrario conditionality loss

The proposed *a contrario* cGAN approach is based on training with unconditionally paired sets of data, obtained by randomly shuffling or re-pairing the original conditional sets of data. The *a contrario* set is defined as $\mathcal{C}_U(\tilde{\mathbf{x}}, \mathbf{y})$, where $\tilde{\mathbf{x}} \sim p(\mathbf{x})$ is the *a contrario* conditional variable ($\tilde{\mathbf{x}} \neq \mathbf{x}$) and \mathbf{y} is the real training variable as in Section 2.1. In this case $\tilde{\mathbf{x}}$ and \mathbf{y} are independent. The generator of the *a contrario* cGAN outputs a transformed set of data $\mathcal{C}_{UG}(\tilde{\mathbf{x}}, \mathbf{y})$ composed of the generator output variable $\mathbf{y}_G \sim p_G(\mathbf{y})$ and the random variable $\tilde{\mathbf{x}}$. For the purpose of this paper these two sets of data will be called "real-*a contrario*" and "generated-*a contrario*" respectively. The motivation to create these new sets is to train the discriminator to correctly classify unconditional data as false. Figure 1 shows the four possible pairings. In practice, random sampling of *a contrario* pairs is carried out without replacement and attention is paid to not include any conditional variables into a same batch while processing.

In order to enforce conditionality between \mathbf{y} and \mathbf{x} an *a contrario* term is proposed as:

$$\mathcal{L}_{ac} = \max_D \left(\mathbb{E}_{\tilde{\mathbf{x}} \sim p(\tilde{\mathbf{x}}), \mathbf{y} \sim p(\mathbf{y})} [\log(1 - D(\tilde{\mathbf{x}}, \mathbf{y}))] + \mathbb{E}_{\tilde{\mathbf{x}} \sim p(\tilde{\mathbf{x}}), \mathbf{x} \sim p(\mathbf{x})} [\log(1 - D(\tilde{\mathbf{x}}, G(\mathbf{x})))] \right) \quad (5)$$

The first term enforces the real-*a contrario* pairs to be classified as fakes. The second term enforces the generated-*a contrario* as fake. The final loss is:

$$\mathcal{L}'_{adv} = \mathcal{L}_{adv} + \mathcal{L}_{ac} \quad (6)$$

3 Experimental section

Several experiments will be presented that evaluate the conditionality of cGANs including: Real image generation from semantic masks on Cityscapes dataset [7]; "Single label"-to-image on CIFAR-10 [28]; Monocular depth estimation on [37]; Semantic segmentation using

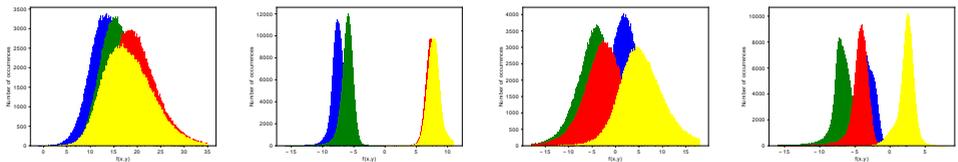


Figure 2: Label-to-image histogram results when validating 500 Cityscape images on a discriminator trained until epoch 200. Blue is Generated-conditional, Green is generated *a contrario*, Red is real-*a contrario*, Yellow is Real-conditional. (a) The trained baseline discriminator, (b) Optimal baseline discriminator, (c) *a contrario* cGAN discriminator, (d) Optimal *a contrario* cGAN discriminator. (a) and (c) are still learning indicating no vanishing gradient or mode collapse [1]. (b) doesn't detect conditionality since *a contrario* real is classified as true (red) (d) succeeds to classify all modes correctly.

pix2pix on Cityscapes dataset. For label-to-image generation, pix2pix[22], pix2pixHD[54], SPADE[41] and CC-FPSE[33] were used to test the conditionality and to highlight the contribution of the *a contrario* cGAN wrt state-of-the-art approaches. "Single label"-to-image is also considered as most new techniques that improves cGANs are designed tested using this task. To demonstrate the generality of the proposed approach depth estimation and image-to-label segmentation tasks are also performed. These structured prediction problems offer strong metrics for evaluating cGANs and various public datasets are available for training. While the scope of conditional evaluation has been limited to tasks that could provide a metric to evaluate both the conditionality and the quality of the generation, the proposed approach is general and not specific to these particular tasks. During training, the network's architecture, the additional losses, the hyper-parameters and data augmentation schemes are kept as in the original papers [22, 33, 41, 54]. The new additional *a contrario* term is the only difference between the compared methods.

3.1 Evaluating conditionality

Preliminary conditionality evaluation follows the method presented in Section 2.2 using *a contrario* sets to evaluate an optimal discriminator. In a first part experiments were carried out on the vanilla pix2pix cGAN with a discriminator PatchGAN architecture with 70×70 receptive field. The model was trained on the Cityscapes dataset [7] for label-to-image translation with 2975 training images resized to 256×256 . The pix2pix GAN model is trained with the same hyper-parameters as specified in the original paper [22]. The evaluation histogram is calculated on the values of the last convolution layer of the discriminator ($f(x, y)$ of Eq (2)) based on the 500 validation images. Each sample from the last convolutional layer is composed of a 30×30 overlapping patches with one channel. The proposed approach is trained in exactly the same manner with the only difference being the new objective function.

Various tests were carried out to investigate the output distributions of each set of data for both the baseline architecture and the proposed method. The underlying accuracy of the implementation was first validated to ensure the accuracy reported in the original paper. A histogram analysis was then performed for different levels of training including: training for 20, 100, 200 epochs and evaluating after each. In another experiment the discriminator was allowed to continue to converge for one epoch after fixing 20, 100 and 200 epochs of cGAN training. In particular, training is performed with the objective given in Eq (4) and as proposed in [1, 12]. These results are plotted for each data pairing: real-conditional, generated-conditional, real-*a-contrario* and generated-*a-contrario* in Figure 2. Figure 2 (a) and (c) show that during training the generator is still learning with no vanishing gradient or mode collapse. In Figure 2 (b) the discriminator has been allowed to reach an optimal value by fixing the generator. The real *a contrario* pairing is wrongly classified 99.9% of the

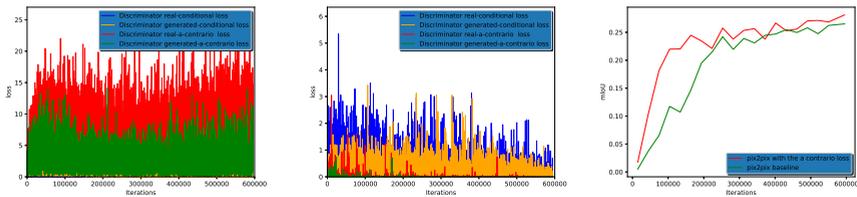


Figure 3: Comparison of the proposed approach on the Cityscape label-to-image training set. (a) The loss function for each set of data-pairing for the baseline cGAN method (*a contrario* are for evaluation only). (b) The loss function for each set of data-pairing for the *a contrario* cGAN method. (c) The evolution of the mIoU for both methods, performed on the validation dataset. It can be seen in (b) compared to (a) that the *a contrario* loss converges to 0 rapidly for the proposed approach. In (c) the proposed approach is much more efficient and converges much faster and with higher accuracy.

time indicating that the discriminator has not learnt conditionality. (d) Shows clearly four distinct distributions and shows the ability of the proposed approach to learn conditionality and correctly classify real *a contrario* pairing 91.9% of the time. Refer to Section A.1 for more insight into the histogram evaluation of conditionality. Similar conditionality tests were performed for various alternative architectures including using a separate/shared network for \mathbf{x} and \mathbf{y} and early/late/at-each-layer fusion. In all cases conditionality was not learnt.

The same analysis was performed for the monocular depth prediction dataset and similar observations were made (see Figure A.1). These results strongly suggest that classic cGAN is unable to learn conditionality and that the spectacular results obtained by cGAN architectures are largely due to higher a level style constraints that are not specific to the input condition variable since swapping condition variables produces no effect. The proposed histogram test allows to demonstrate the ability of the discriminator to classify the various underlying classes of data and shows their statistical distribution. The approach proposed in [1, 26] looks only at accuracy, while the proposed histogram test also allows to see how well the discriminator distinguishes between the different data modalities (real/fake/*a contrario*).

3.2 Label-To-Image translation

Generating realistic images from semantic labels is good task to evaluate the effect of the *a contrario* at a high level since many images can be potentially generated for each semantic class label. Figure 3(c) shows a comparison of the mIoU for the baseline pix2pix model and proposed pix2pix model with the additional *a contrario* loss. It can be observed that the *a contrario* cGAN converges faster than the baseline. The mIoU of the model with *a contrario* at iteration 163k is 24.46 whereas the baseline is 14.65. The mIoU oscillates around that value for the *a contrario* model indicating that the model has converged. After 595k iterations, the mIoU for both models are very close 28.28 and 26.41. It is worth noting that evaluating using real images yields 29.6. The convergence is reported for the generator where the computational cost is exactly the same. the *a contrario* loss is specific only to the discriminator and adds a small computational cost. By restricting the search space of the generator to only conditional pairs, the generator’s convergence is faster.

Table 1 shows a comparison of different architectures with and without *a contrario* augmentation. For a fair comparison all the networks are trained from scratch and the same hyper-parameter are used. The *a contrario* loss is the only difference between the two networks. The batch size for SPADE is 32 and 16 for CC-FPSE. Through explicitly enforcing the conditionality with *a contrario* examples, the discriminator learns to penalize unconditional generation achieving better results. Figure A.10 shows a qualitative comparison.

Moreover, Figure 3(a) and Figure 3(b) show the comparison of the losses of the discrim-

Method	Resolution	FID	mIoU	Pixel accuracy (PA)
pix2PixHD	256 × 512	66.7	56.9	92.8
<i>a contrario</i> pix2pixHD	256 × 512	60.1	60.1	93.2
SPADE	256 × 512	65.5	60.2	93.1
<i>a contrario</i> SPADE	256 × 512	59.9	61.5	93.7
CC-FPSE	256 × 512	52.4	61.8	92.8
<i>a contrario</i> CC-FPSE	256 × 512	53.5	63.9	93.5

Table 1: A comparison of different architectures trained from scratch with and without a *a contrario* augmentation. The networks with *a contrario* achieves better results with a mean improvement of $\Delta mIoU = +2.3$, $\Delta PA = +0.56$, and $\Delta FID = -3.8$.

inator for both models on this dataset. The baseline is trained with only conditional pairs, however, the *a contrario* data pairs are plotted to asses the ability of the discriminator to learn the conditionality automatically. The *a contrario* losses remain high for the baseline and converge to 0 for the proposed *a contrario* cGAN. Figure 2 presented the histogram results for this experiment showed that the proposed approach better models conditionality. Subsequently the gradients provided to the generator are expected to be better directed. The gradient norm is plotted in Figure A.5.

3.3 "Single-label"-to-image

The generality of the proposed *a contrario* cGAN can also be demonstrated by showing that it also improves architectures other than image-to-image. An example of a different task is conditioning the generated image on a single input class-label as in [4, 24, 25, 36, 40]. This different architecture is of interest because many new methods for improving cGANs are often tested on this task. Unfortunately, these methods are mainly evaluated on the FID [18] and IS [45] scores. As stated earlier, these metrics measure the quality/diversity and they favor models that memorise the training set [15]. They have not been designed to evaluate conditionality and therefore not sufficient for the purpose of this paper. Despite that, these criteria are still important for evaluating the quality of GANs, however, an additional criterion is required for testing conditionality.

Here a simple conditionality test is proposed specifically for "single label"-to-image generation tasks based on a pretrained Resnet-56 [16] classifier trained on CIFAR-10 [28]. BigGAN [4] was selected as the baseline. Since BigGAN uses the Hinge-loss [31], the *a contrario* loss is adapted as follows:

$$\begin{aligned}
 \mathcal{L}_D &= -\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}), \mathbf{y} \sim p(\mathbf{y}|\mathbf{x})} [\min(0, -1 + D(\mathbf{x}, \mathbf{y}))] - \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\min(0, -1 - D(\mathbf{x}, G(\mathbf{x})))] \\
 &\quad - \mathbb{E}_{\tilde{\mathbf{x}} \sim p(\tilde{\mathbf{x}}), \mathbf{y} \sim p(\mathbf{y})} [\min(0, -1 - D(\tilde{\mathbf{x}}, \mathbf{y}))] - \mathbb{E}_{\tilde{\mathbf{x}} \sim p(\tilde{\mathbf{x}}), \mathbf{x} \sim p(\mathbf{x})} [\min(0, -1 - D(\tilde{\mathbf{x}}, G(\mathbf{x})))] \\
 \mathcal{L}_G &= -\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} D(\mathbf{x}, G(\mathbf{x}))
 \end{aligned} \tag{7}$$

Both models are trained from scratch on CIFAR-10 [28] dataset using the hyper-parameter specified in [4]. The conditionality is tested by generating 10k images for each label(100k images in total) and calculating the accuracy. The results¹are shown in Table 2.

Method	IS score	FID score	Acc
BigGAN[4]	8.26 ± 0.095	6.84	86.54
<i>a contrario</i> BigGAN	8.40 ± 0.067	6.28	92.04

Table 2: A comparison of BigGAN [4] with and without the *a contrario* GAN. The network with *a contrario* achieves significantly better results with an improvement of $\Delta Acc = +5.59$, $\Delta IS = +0.14$, and $\Delta FID = -0.56$.

¹The Pytorch IS and FID implementations were used for comparison

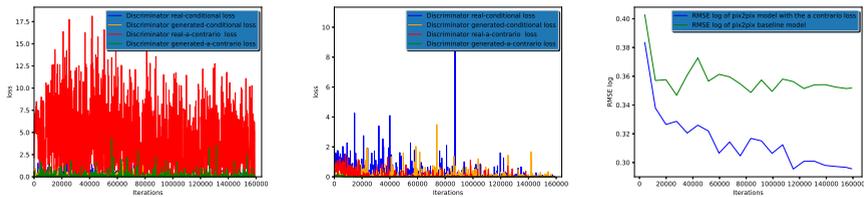


Figure 4: Comparison of the proposed approach on the NYU Depth V2 training set. (a) The loss function for each set of data-pairing for the baseline cGAN method. (b) The loss function for each set of data-pairing for the *a contrario* cGAN method. (c) The evolution of the RMSE log for both methods, performed on the validation dataset. It can be seen in (b) compared to (a) that the *a contrario* loss converges to 0 rapidly for the proposed approach. In (c) the proposed approach is much more efficient and converges much faster and with much higher accuracy.

The conditionality improved significantly over the baseline with $\Delta Acc = +5.59$ and the quality also improved with $\Delta FID = -0.56$, $\Delta IS = +0.14$. Similar to the observation made before *a contrario* enforces the conditionality without compromising the quality. A failure mode of the lack of conditionality of the discriminator is class leakage : images from one class contain properties of another. While it is not easy to define a proper metric for such failure mode, it is shown that using the *a contrario* loss the classification was improved and therefore the generation is better constrained and does not mix class properties. This result shows that *a contrario* GAN also improves on a different SOTA task and confirms again that conditionality is an overlooked factor in current SOTA metrics.

3.4 Monocular Depth prediction

Monocular depth prediction is an ill-posed problem as an infinite number of 3D scenes can be projected onto the same 2D scene. A good model for predicting image geometry should exploit visual clues such as object sizes, lighting, shadows, object localization, perspective and texture information. Depth prediction is an appropriate task to evaluate the conditional relationship between \mathbf{x} and \mathbf{y} . The RMSE log error is a strong metric to evaluate the conditional information since it takes into account the fidelity of the prediction with respect to the input. Similar to [11] other metrics are also reported for completeness. The model is trained on the NYU Depth V2 Dataset [37] to predict depth from monocular 2D-RGB images only. The official train/validation split of 795 pairs is used for training and 694 pairs are used for validation. Similarly to low resolution label-to-image translation, the dataset images are resized to have a resolution of 256×256 . The experiment is repeated 6 times without fixing the seed and the mean and standard deviation are reported.

Figure 4(c) shows the RMSE log performances of the validation set over the training iterations for one experiment. Table 3 shows the comparison of the two models across different metrics. Clearly the *a contrario* cGAN reaches a better performance with log RMSE 0.3036 versus 0.3520 for the baseline (the mean is reported here). The evaluation of mode collapse using NDB [44] of the two networks is given in Section A.3. The qualitative results are shown in Figure A.8. Moreover, Figure 4(a) and Figure 4(b) compare the losses of the discriminator for the different data pairings when using the baseline model and the proposed *a contrario* model. Similar to Label-To-Image translation, the baseline results show high loss for real *a contrario*.

The discriminator is optimized only to distinguish real and generated samples its decision boundary is independent off conditional variable. The baseline cGAN architecture will not penalize the generation of outputs belonging to the target domain but that do not correspond to the input. Not only does this leave the generator with a larger search space (the generator

Method	RMSE log	silog	\log_{10}	abs rel
baseline	0.3520 \pm 0.0016	28.54 \pm 0.1932	0.1247 \pm 0.0005	0.3318 \pm 0.0026
<i>a contrario</i>	0.3036 \pm 0.0055	23.51 \pm 0.1932	0.1079 \pm 0.0021	0.2868 \pm 0.0093

Table 3: Monocular Depth prediction experiments were repeated on the baseline and *a contrario* cGANs 6 times with different seeds. The mean and standard deviation are reported for each metric. The results shows that the *a contrario* cGAN outperforms the baseline on the depth metric [11].

Method	Pixel accuracy (PA)	Mean Acc	FreqW Acc	mIoU
Baseline	66.12	23.31	53.64	15.97
<i>a contrario</i>	72.93	26.87	60.40	19.23

Table 4: Comparison on the Cityscapes dataset validation set. The proposed method consistently obtains more accurate results and finishes with a largely different score at the end of training with mIoU of 19.23 versus for the baseline 15.97.

is less efficient), but it can allow mode collapse whereby the generator always produces the same output. The *a contrario* loss explicitly avoids this by penalizing unconditional generation. The *a contrario* losses of the *a contrario* cGAN can be observed in Figure 4(a) and Figure 4(b) to converge rapidly to 0. This indicates the search space of the generator is explicitly restricted to only conditional samples early in the training.

3.5 Image-to-label segmentation

Image-to-label is a simpler task compared to depth prediction and label-to-image prediction as the goal of the generator is to transfer from a high-dimensional space to a lower-dimensional space. Furthermore, the evaluation is simpler since the image mask does not have multiple solutions and it is not necessary to use an external pre-trained segmentation network for comparison as in the case of label-to-image translation. It is worth mentioning that pix2pix is trained to output 19 classes as a segmentation network and is not trained as an image-to-image network as it is often done in cGAN architectures. FCN [34] trained on [22] obtains 21.0 mIoU. The performances are shown in Table 4. As expected the training was unstable (refer to Figure A.6). However, the *a contrario* cGAN shows superior mIoU performance with 19.23 versus 15.97 for the baseline model. Figure A.7 shows the qualitative results of the both models. It can be observed that the model baseline has invented labels that are not specified by the input. Training with *a contrario* helps the discriminator to model conditionality. Thus, the generator search space is restricted to only conditional space. The generator is penalized for conditionality even if the generation is realistic.

4 Conclusion

This paper has proposed two important contributions for conditional Generative Adversarial Networks (cGANs). The first main contribution is a probabilistic analysis of the discriminator to show that it is not explicitly conditional. The second contribution is a new method, called *a contrario* cGAN, that explicitly models conditionality for both parts of the adversarial architecture via a novel *a contrario* loss that involves training the discriminator to learn unconditional (adverse) examples. This leads to a novel type of data augmentation approach for GANs (*a contrario* learning) which allows to restrict the search space of the generator to conditional outputs using adverse examples. Extensive experimentation has shown significant improvements across several tasks, datasets and architectures.

Although *a contrario* has shown significant improvement for image domain addressing other modalities such as audio and text is to be tested in the future work. It will also be of interest to explore multiple discriminator output classes for the four pairings considered in this paper.

5 Acknowledgements

This work was performed using HPC resources from GENCI-IDRIS (Grant 2021-011011931). The authors would like to acknowledge the Association Nationale Recherche Technologie (ANRT) for CIFRE funding (n°2019/1649).

References

- [1] Martín Arjovsky and L. Bottou. Towards principled methods for training generative adversarial networks. *ArXiv*, abs/1701.04862, 2017.
- [2] Sanjeev Arora, R Ge, Y Liang, T Ma, and Y Zhang. Generalization and equilibrium in generative adversarial nets (gans). In *34th International Conference on Machine Learning, ICML 2017*, 2017.
- [3] Sanjeev Arora, Andrej Risteski, and Yi Zhang. Do gans learn the distribution? some theory and empirics. In *International Conference on Learning Representations*, 2018.
- [4] A. Brock, J. Donahue, and K. Simonyan. Large scale gan training for high fidelity natural image synthesis. *ArXiv*, abs/1809.11096, 2019.
- [5] C. Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A. Efros. Everybody dance now. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5932–5941, 2019.
- [6] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1511–1520, 2017.
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [8] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 113–123, 2019.
- [9] Zihang Dai, Zhilin Yang, Fan Yang, William W Cohen, and Ruslan Salakhutdinov. Good semi-supervised learning that requires a bad gan. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6513–6523, 2017.
- [10] Vincent Dumoulin, Jonathon Shlens, and M. Kudlur. A learned representation for artistic style. *ArXiv*, abs/1610.07629, 2017.
- [11] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in Neural Information Processing Systems*, 27:2366–2374, 2014.
- [12] Farzan Farnia and A. Ozdaglar. Gans may have no nash equilibria. *ArXiv*, abs/2002.09124, 2020.

- [13] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. *arXiv preprint arXiv:2012.07177*, 2020.
- [14] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems-Volume 2*, pages 2672–2680, 2014.
- [15] Ishaan Gulrajani, Colin Raffel, and Luke Metz. Towards gan benchmarks which require generalization. *arXiv preprint arXiv:2001.03653*, 2020.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [17] Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *International Conference on Learning Representations*, 2019.
- [18] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, 2017.
- [19] Gabriel Huang, Hugo Berard, Ahmed Touati, Gauthier Gidel, Pascal Vincent, and Simon Lacoste-Julien. Parametric adversarial divergences are good task losses for generative modeling. *arXiv preprint arXiv:1708.02511*, 2017.
- [20] X. Huang and Serge J. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1510–1519, 2017.
- [21] Daniel Jiwoong Im, He Ma, Graham W Taylor, and Kristin Branson. Quantitatively evaluating gans with divergences proposed for training. In *International Conference on Learning Representations*, 2018.
- [22] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5967–5976, 2017.
- [23] J. Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016.
- [24] Minguk Kang and Jaesik Park. Contragan: Contrastive learning for conditional image generation. *arXiv: Computer Vision and Pattern Recognition*, 2020.
- [25] Tero Karras, Miika Aittala, Janne Hellsten, S. Laine, J. Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *ArXiv*, abs/2006.06676, 2020.
- [26] Ilya Kavalerov, Wojciech Czaja, and Rama Chellappa. A multi-class hinge loss for conditional gans. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1290–1299, 2021.
- [27] Naveen Kodali, James Hays, J. Abernethy, and Z. Kira. On convergence and stability of gans. *arXiv: Artificial Intelligence*, 2018.

- [28] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research).
- [29] Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiri Matas. Deblurgan: Blind motion deblurring using conditional adversarial networks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8183–8192, 2018.
- [30] Wenbo Li, Pengchuan Zhang, Lei Zhang, Qiuyuan Huang, X. He, Siwei Lyu, and Jianfeng Gao. Object-driven text-to-image synthesis via adversarial training. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12166–12174, 2019.
- [31] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017.
- [32] Lingjie Liu, Weipeng Xu, M. Zollhöfer, H. Kim, F. Bernard, Marc Habermann, W. Wang, and C. Theobalt. Neural rendering and reenactment of human actor videos. *ACM Transactions on Graphics (TOG)*, 38:1 – 14, 2019.
- [33] Xihui Liu, Guojun Yin, Jing Shao, Xiaogang Wang, and hongsheng Li. Learning to predict layout-to-image conditional convolutions for semantic image synthesis. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [34] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [35] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [36] Takeru Miyato and Masanori Koyama. cgans with projection discriminator. *ArXiv*, abs/1802.05637, 2018.
- [37] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012.
- [38] Weili Nie and A. Patel. Jr-gan: Jacobian regularization for generative adversarial networks. *ArXiv*, abs/1806.09235, 2018.
- [39] Evangelos Ntavelis, A. Romero, I. Kastanis, L. Gool, and R. Timofte. Sesame: Semantic editing of scenes by adding, manipulating or erasing objects. *ArXiv*, abs/2004.04977, 2020.
- [40] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *ICML*, 2017.
- [41] T. Park, Ming-Yu Liu, T. Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2332–2341, 2019.
- [42] Y. Qu, Yizi Chen, J. Huang, and Yuan Xie. Enhanced pix2pix dehazing network. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8152–8160, 2019.

- [43] S. Reed, Zeynep Akata, Xinchun Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text to image synthesis. In *ICML*, 2016.
- [44] Eitan Richardson and Yair Weiss. On gans and gmms. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 5852–5863, 2018.
- [45] Tim Salimans, Ian J. Goodfellow, W. Zaremba, Vicki Cheung, A. Radford, and Xi Chen. Improved techniques for training gans. In *NIPS*, 2016.
- [46] Abhishek Sinha, Kumar Ayush, Jiaming Song, Burak Uzkent, Hongxia Jin, and Stefano Ermon. Negative data augmentation. In *International Conference on Learning Representations*, 2021.
- [47] A. Srivastava, L. Valkov, Chris Russell, Michael U. Gutmann, and Charles Sutton. Veegan: Reducing mode collapse in gans using implicit variational learning. In *NIPS*, 2017.
- [48] Vadim Sushko, Edgar Schönfeld, D. Zhang, Juergen Gall, B. Schiele, and A. Khoreva. You only need adversarial supervision for semantic image synthesis. *ArXiv*, abs/2012.04781, 2020.
- [49] Hao Tang, D. Xu, Yan Yan, P. H. S. Torr, and N. Sebe. Local class-specific and global image-level generative adversarial networks for semantic-guided scene generation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7867–7876, 2020.
- [50] N. T. Tran, V. H. Tran, N. B. Nguyen, T. K. Nguyen, and N. M. Cheung. On data augmentation for gan training. *IEEE Transactions on Image Processing*, 30:1882–1897, 2021. doi: 10.1109/TIP.2021.3049346.
- [51] Thomas Unterthiner, Bernhard Nessler, G. Klambauer, Martin Heusel, Hubert Ramsauer, and S. Hochreiter. Coulomb gans: Provably optimal nash equilibria via potential fields. *ArXiv*, abs/1708.08819, 2018.
- [52] H. D. Vries, Florian Strub, J. Mary, H. Larochelle, Olivier Pietquin, and Aaron C. Courville. Modulating early visual processing by language. In *NIPS*, 2017.
- [53] T. Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, A. Tao, J. Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *NeurIPS*, 2018.
- [54] T. Wang, Ming-Yu Liu, Jun-Yan Zhu, A. Tao, J. Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8798–8807, 2018.
- [55] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. Ieee, 2003.
- [56] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [57] T. Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and X. He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1316–1324, 2018.

- [58] H. Zhang, Vishwanath Sindagi, and V. Patel. Image de-raining using a conditional generative adversarial network. *IEEE Transactions on Circuits and Systems for Video Technology*, 30:3943–3956, 2020.
- [59] H. Zhang, Zizhao Zhang, Augustus Odena, and H. Lee. Consistency regularization for generative adversarial networks. *ArXiv*, abs/1910.12027, 2020.
- [60] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1947–1962, 2018.
- [61] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [62] Xinyu Zhang, Qiang Wang, Jian Zhang, and Zhao Zhong. Adversarial autoaugment. In *International Conference on Learning Representations*, 2019.
- [63] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. *arXiv preprint arXiv:2006.10738*, 2020.
- [64] Zhengli Zhao, Sameer Singh, H. Lee, Zizhao Zhang, Augustus Odena, and Han Zhang. Improved consistency regularization for gans. *ArXiv*, abs/2002.04724, 2020.
- [65] Zhengli Zhao, Zizhao Zhang, Ting Chen, Sameer Singh, and Han Zhang. Image augmentations for gan training. *arXiv preprint arXiv:2006.02595*, 2020.
- [66] P. Zhou, Lingxi Xie, Xiaopeng Zhang, B. Ni, and Q. Tian. Searching towards class-aware generators for conditional generative adversarial networks. *ArXiv*, abs/2006.14208, 2020.
- [67] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.