

This CVPR 2021 workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

One-Shot GAN: Learning to Generate Samples from Single Images and Videos

Figure 1: Our proposed One-Shot GAN needs only one video (first two rows) or one image (last two rows) for training. At inference phase, it generates novel scene compositions with varying content and layouts. E.g., from a single video with a car on a road, One-Shot GAN can generate the scene without the car or with two cars, and for a single air balloon image, it produces layouts with different number and placement of the balloons. (Original samples are shown in grey or red frames.)

Abstract

Training GANs in low-data regimes remains a challenge, as overfitting often leads to memorization or training divergence. In this work, we introduce One-Shot GAN that can learn to generate samples from a training set as little as one image or one video. We propose a two-branch discriminator, with content and layout branches designed to judge the internal content separately from the scene layout realism. This allows synthesis of visually plausible, novel compositions of a scene, with varying content and layout, while preserving the context of the original sample. Compared to previous single-image GAN models, One-Shot GAN achieves higher diversity and quality of synthesis. It is also not restricted to the single image setting, successfully learning in the introduced setting of a single video.

1. Introduction

Without sufficient training data, GANs are prone to overfitting, which often leads to mode collapse and training instabilities [9, 4]. This dependency on availability of training data limits the applicability of GANs in domains where collecting a large dataset is not feasible. In some real-world applications, collection even of a small dataset remains challenging. It may happen that rare objects or events are present only in one image or in one video, and it is difficult to obtain a second one. This, for example, includes pictures of exclusive artworks or videos of traffic accidents recorded in extreme conditions. Enabling learning of GANs in such *one-shot* scenarios has thus a potential to improve their utilization in practice. Previous work [9, 4] studied one-shot image generation in the context of learning from a *single image*. In this work, we introduce a novel setup of learn-



Figure 2: One-Shot GAN. The two-branch discriminator judges the content distribution separately from the scene layout realism and thus enables the generator to produce images with varying content and global layouts. See Sec. 2 for details.

ing to generate new images from frames of a *single video*. In practice, recoding a video lasting for several seconds can take almost as little effort as collecting one image. However, a video contains much more information about the scene and the objects of interest (e.g., different poses and locations of objects, various camera views). Learning from a video can enable generation of images of higher quality and diversity, while still operating in a one-shot mode, and therefore can improve its usability for applications.

To mitigate overfitting in the one-shot mode, recent single-image GAN models [9, 4] proposed to learn an image patch-based distribution at different image scales. Though these models overcome memorization, producing different versions of a training image, they cannot learn high-level semantic properties of the scene. They thus often suffer from incoherent shuffling of image patches, distorting objects and producing unrealistic layouts (see Fig. 3 and 4). Other low-data GAN models, such as FastGAN [6], have problems with memorization when applied in the one-shot setting (see Sec. 3). In this work, we go beyond patch-based learning, seeking to generate novel plausible compositions of objects in the scene, while preserving the original image context. Thus, we aim to keep novel compositions visually plausible, with objects preserving their appearance, and the scene layout looking realistic to a human eye.

To this end, we introduce One-Shot GAN, an unconditional single-stage GAN model, which can generate images that are significantly different from the original training sample while preserving its context. This is achieved by two key ingredients: the novel design of the discriminator and the proposed diversity regularization technique for the generator. The new One-Shot GAN discriminator has two branches, responsible for judging the content distribution and the scene layout realism of images separately from each other. Disentangling the discriminator decision about the content and layout helps to prevent overfitting and provides more informative signal to the generator. To achieve high diversity of generated samples, we also extend the regularization technique of [13, 2] to one-shot unconditional image synthesis. As we show in Sec. 3, The proposed One-Shot GAN generates high-quality images that are significantly different from training data. One-Shot GAN is the first model that succeeds in learning from both single images and videos, improving over prior work [9, 6] in image quality and diversity in these one-shot regimes.

2. One-Shot GAN

Content-Layout Discriminator. We introduce a solution to overcome the memorization effect but still to generate images of high quality in the one-shot setting. Building on the assumption that to produce realistic and diverse images the generator should learn the appearance of objects and combine them in a globally-coherent way in an image, we propose a discriminator judging the *content* distribution of an image separately from its layout realism. To achieve the disentanglement, we design a two-branch discriminator architecture, with separate content and layout branches (see Fig.2). Our discriminator D consists of the low-level feature extractor $D_{low-level}$, the content branch $D_{content}$, and the layout branch D_{layout} . Note that the branching happens after an intermediate layer in order to learn a relevant representation. $D_{content}$ judges the content of this representation irrespective from its spatial layout, while D_{layout} , inspects only the spatial information. Inspired by the attention modules of [7, 12], we extract the *content* from intermediate representations by aggregating spatial information via global average pooling, and obtain *layout* by aggregating channels via a simple (1×1) convolution. This way, the content branch judges the fidelity of objects composing the image independent of their spatial location, while the layout branch is sensitive only to the realism of the global scene layout. Note that $D_{content}$ and D_{layout} receive only limited information from previous layers, which prevents overfitting. This helps to overcome the memorization of training data and to produce different images.

Diversity regularization. To improve variability of generated samples, we propose to add diversity regularization (DR) loss term \mathcal{L}_{DR} to the objective. Previously proposed



Figure 3: Comparison with other methods in the single image setting. Single-image GAN of [9] is prone to incoherently shuffle image patches (e.g. sky textures appear on the ground), and the few-shot FastGAN model [6] collapses to producing the original image or its flipped version. In contrast, One-Shot GAN produces diverse images with realistic *global* layouts.

	Single Image				Single Video			
Method	SIFID↓	LPIPS↑	$\left \text{ MS-SSIM} \right\downarrow$	Dist.	SIFID↓	LPIPS↑	MS-SSIM↓	Dist.
				to train				to train
SinGAN [9]	0.13	0.26	0.69	0.30	2.47	0.32	0.65	0.51
FastGAN [6]	0.13	0.18	0.77	0.11	0.79	0.43	0.55	0.13
One-Shot GAN	0.08	0.33	0.63	0.37	0.55	0.43	0.54	0.34

Table 1: Comparison in the Single Image and Single Video settings on DAVIS-YFCC100M [8, 10] dataset.

regularization terms [13, 2] aimed to encourage the generator to produce different outputs depending on the input latent code, in such a way that the generated samples with closer latent codes should look more similar to each other, and vice versa. In contrast, in the one-shot image synthesis setting, the perceptual distance of the generated images should not be dependent on the distance between their latent codes. As we operate in one semantic domain, the generator should produce images that are in-domain but more or less equally different from each other and substantially different from the original sample. Thus, we propose to encourage the generator to produce perceptually different image samples independent of their distance in the latent space. \mathcal{L}_{DR} is expressed as follows:

$$\mathcal{L}_{DR}(G) = \mathbb{E}_{z_1, z_2} \Big[\frac{1}{L} \sum_{l=1}^{L} \| G^l(z_1) - G^l(z_2) \| \Big], \quad (1)$$

where $\|\cdot\|$ denotes the L1 norm, $G^l(z)$ indicates features extracted from the *l*-th block of the generator G given the latent code z. Contrary to prior work, we compute the distance between samples in the feature space, enabling more meaningful diversity of the generated images, as different generator layers capture various image semantics, inducing both high- and low-level diversity.

Final objective. We compute adversarial loss for each discriminator part: $D_{low-level}$, $D_{content}$, and D_{layout} . This way, the discriminator decision is based on low-level details of images, such as textures, and high-scale properties, such as content and layout. The overall adversarial loss is

$$\mathcal{L}_{adv}(G,D) = \mathcal{L}_{D_{content}} + \mathcal{L}_{D_{layout}} + 2\mathcal{L}_{D_{low-level}}, \quad (2)$$

where \mathcal{L}_{D_*} is the binary cross-entropy $\mathbb{E}_x[\log D_*(x)] + \mathbb{E}_x[\log(1 - D_*(G(z)))]$ for real image x and generated im-

age G(z). As the two branches of the discriminator operate at high-level image features, contrary to only one $D_{low-level}$ operating at low-level features, we double the weighting for the $\mathcal{L}_{D_{low-level}}$ loss term. This is done in order to properly balance the contributions of different feature scales and encourage the generation of images with good low-level details, coherent contents and layouts.

The overall One-Shot GAN objective can be written as:

$$\min_{G} \max_{D} \mathcal{L}_{adv}(G, D) - \lambda \mathcal{L}_{DR}(G), \qquad (3)$$

where λ controls the strength of the diversity regularization and \mathcal{L}_{adv} is the adversarial loss in Eq. 2.

Implementation. The One-Shot GAN discriminator uses ResNet blocks, following [1]. We use three ResNet blocks before branching and four blocks for the content and layout branches. We employ standard image augmentation strategies for the discriminator training, following [5]. λ for \mathcal{L}_{DR} in Eq. 3 is set to 0.15. We use the ADAM optimizer with $(\beta_1, \beta_2) = (0.5, 0.999)$, a batch size of 5 and a learning rate of 0.0002 for both G and D.

3. Experiments

Evaluation settings. We evaluate One-Shot GAN on two different one-shot settings: training on a single image and a single video. We select 15 videos from DAVIS [8] and YFCC100M [10] datasets. In the Single Video setting, we use all frames of a video as training images, while for the Single Image setup we use only one middle frame. The chosen videos last for 2-10 seconds and consist of 60-100 frames. To assess the quality of generated images, we measure single FID (SIFID) [9]. Image diversity is assessed by the average LPIPS [3] and MS-SSIM [11] across pairs of



Figure 4: Comparison with other methods in the single video setting. While other models fall into reproducing the training frames or fail to learn textures, One-Shot GAN generates high-quality images significantly different from the original video.

generated images. To verify that the models do not simply reproduce the training set, we report average LPIPS to the nearest image in the training set, augmented in the same way as during training (Dist. to train). We compare our model with a single image method SinGAN [9] and with a recent model on few-shot image synthesis, FastGAN [6].

Main results. Table 1 presents quantitative comparison between the models in the Single Image and Video settings, while the respective visual results are shown in Fig. 3 and 4. As seen from Table 1, One-Shot GAN notably outperforms other models in both quality and diversity metrics. Importantly, our model is the only one which successfully learns from both single images and single videos.

As seen from Fig. 1 and 3, in the Single Image setting, One-Shot GAN produces diverse samples of high visual quality. For example, our model can change the number of rocks on the background or change their shapes. Note that such changes keep appearance of objects, preserving original content, and maintain scene layout realism. In contrast, single-image method SinGAN disrespects layouts (e.g. sky textures may appear below horizon), and is prone to modest diversity, especially around image corners. This is reflected in higher SIFID and lower diversity in Table 1. The fewshot FastGAN suffers from memorization, only reproducing the training image or its flipped version. In Table 1 this is reflected in low diversity and small Dist. to train (in red).

In the proposed Single Video setting, there is much more information to learn from, so generative models can learn more interesting combinations of objects and scenes. Fig. 1 and 4 show images generated by the models in this setting. One-Shot GAN produces plausible images that are substantially different from the training frames, adding/removing objects and changing scene geometry. For example, having seen a bus following a road, One-Shot GAN varies the length of a bus and placement of trees. For the video with an equestrian competition, our model can remove a horse from the scene and change the jumping obstacle configuration. In contrast, SinGAN, which is tuned to learn from a single image, does not generalize to this setting, producing "mean in class" textures and failing to learn appearance of objects (low diversity and very high SIFID). FastGAN, on the other hand, learns high-scale scene properties, but fails to augment the training set with non-trivial changes, having a very low distance to the training data (0.13 in Table 1).

Table 1 confirms that the proposed two-branch discriminator in combination with diversity regularization manages to overcome the memorization effect, achieving high distance to training data in both settings (0.37 and 0.34). This means that One-Shot GAN augments the training set with structural transformations that are orthogonal to standard data augmentation techniques, such as horizontal flipping or color jittering. To achieve this, the model requires as little data as one image or one short video clip. We believe, such ability can be especially useful to generate samples for augmentation of limited data, for example by creating new versions of rare examples.

4. Conclusion

We propose One-Shot GAN, a new unconditional generative model operating at different one-shot settings, such as learning from a single image or a single video. At such lowdata regimes, our model mitigates the memorization problem and generates diverse images that are structurally different from the training set. Particularly, our model is capable of synthesizing images with novel views and different positions of objects, preserving their visual appearance. We believe, such structural diversity provides a useful tool for image editing applications, as well as for data augmentation in domains, where data collection remains challenging.

References

- Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations (ICLR)*, 2019. 3
- [2] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. 2, 3
- [3] Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, Advances in Neural Information Processing Systems (NeurIPs), 2016. 3
- [4] Tobias Hinz, Matthew Fisher, Oliver Wang, and Stefan Wermter. Improved techniques for training single-image gans. arXiv preprint arXiv:2003.11512, 2020. 1, 2
- [5] Tero Karras, Miika Aittala, Janne Hellsten, S. Laine, J. Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In Advances in Neural Information Processing Systems (NeurIPS), 2020. 3
- [6] Bingchen Liu, Yizhe Zhu, Kunpeng Song, and Ahmed Elgammal. Towards faster and stabilized {gan} training for high-fidelity few-shot image synthesis. In *International Conference on Learning Representations*, 2021. 2, 3, 4
- [7] Jongchan Park, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Bam: Bottleneck attention module. *arXiv preprint arXiv:1807.06514*, 2018. 2
- [8] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017. 3
- [9] Tamar Rott Shaham, Tali Dekel, and T. Michaeli. Singan: Learning a generative model from a single natural image. In *International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 3, 4
- [10] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016. 3
- [11] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In Asilomar Conference on Signals, Systems & Computers, 2003. 3
- [12] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision* (ECCV), pages 3–19, 2018. 2
- [13] Dingdong Yang, Seunghoon Hong, Y. Jang, T. Zhao, and H. Lee. Diversity-sensitive conditional generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2019. 2, 3