

---

# Ways of Conditioning Generative Adversarial Networks

---

**Hanock Kwak and Byoung-Tak Zhang**  
School of Computer Science and Engineering  
Seoul National University  
Seoul 151-744, Korea  
{hnikwak, btzhang}@bi.snu.ac.kr

## Abstract

The GANs are generative models whose random samples realistically reflect natural images. It also can generate samples with specific attributes by concatenating a condition vector into the input, yet research on this field is not well studied. We propose novel methods of conditioning generative adversarial networks (GANs) that achieve state-of-the-art results on MNIST and CIFAR-10. We mainly introduce two models: an information retrieving model that extracts conditional information from the samples, and a spatial bilinear pooling model that forms bilinear features derived from the spatial cross product of an image and a condition vector. These methods significantly enhance log-likelihood of test data under the conditional distributions compared to the methods of concatenation.

## 1 Introduction

The goal of the generative model is to learn the underlying probabilistic distribution of unlabeled data by disentangling explanatory factors in the data[1, 2]. It is applicable to tasks such as classification, regression, visualization, and policy learning in reinforcement learning. The GAN[3] is a prominent generative model that is found to be useful for realistic sample generation[4], semi-supervised learning[5], super resolution[6], text-to-image[7], and image inpainting[8]. The potentiality of GAN was supported by the theory that if the model has enough capacity, the learned distribution can converge to the distribution over real data[3]. The representative power of the GAN was highly enhanced with deep learning techniques[4] and various methods[5, 9] was introduced to stabilize the learning process.

After the conditional GAN[10] was first introduced to generate samples of specific labels from a single generator, further research was not made thoroughly despite of its practical usage. For example, it was used to generate images from descriptive sentences[7] or attribute vectors[11]. The conditioned GANs learn conditional probability distribution in which the condition can be any kind of auxiliary information describing the data. Mostly it was modeled by concatenating condition vectors into some layers of the generator and discriminator in GAN. Even though this forms a joint representation, it is hard to fully capture the complex associations between the two different modalities[12]. We introduce a small variant of bilinear pooling that provides multiplicative interaction between all elements of two vectors. To provide computationally efficient and spatially sensible bilinear operation on an image (order 3 tensor) and a vector, we calculated cross product over last dimension of the image, resulting in a new image with increased channels.

If there is an oracle that can extract conditional information perfectly from any sample, we can train GAN to generate samples that the oracle can figure out given conditions from them. This is an information retrieval model where there is a pre-trained model that plays the role of an oracle. In addition to the objective of GAN, it maximizes lower bound of the mutual information between a

given condition and the extracted condition. This idea has originated from infoGAN[13] which is an information-theoretic extension to the GAN that is able to learn disentangled representations in a completely unsupervised manner.

## 2 Related Works

The conditional GAN[10] concatenates condition vector into the input of the generator and the discriminator. Variants of this method was successfully applied in [7, 11, 14]. [7] obtained visually-discriminative vector representation of text descriptions and then concatenated that vector into every layer of the discriminator and the noise vector of the generator. [11] used a similar method to generate face images from binary attribute vectors such as hair styles, face shapes, etc. In [14], Structure-GAN generates surface normal maps and then they are concatenated into noise vector of Style-GAN to put styles in those maps.

The spatial bilinear pooling was mainly inspired from studies on multimodal learning[15]. The key question of the multimodal learning is how can a model uncover the correlated interaction of two vectors from different domains. In order to achieve this, various methods (vector concatenation[16], element-wise operations[17], factorized restricted Boltzmann machine (RBM)[15], bilinear pooling[12, 18], etc) were proposed for numerous challenging tasks. The RBM based models require expensive MCMC sampling which makes it difficult to scale them to large datasets. The bilinear pooling is more expressive than vector concatenation or element-wise operations, but they are inefficient due to squared complexity  $O(n^2)$ . To solve this problem, [19] addressed the space and time complexity of bilinear features using Tensor Sketch[20].

The information retrieving model uses core algorithm of infoGAN[13] that recover disentangled representations by maximizing the mutual information for inducing latent codes. In infoGAN, the input noise vector is decomposed into a source of incompressible noise and the latent code, and there is an auxiliary output in the discriminator to retrieve the latent codes. The infoGAN utilizes Variational Information Maximization[21] to deal with intractability of the mutual information. Unlike infoGAN which randomly generates latent codes, we explicitly put condition information in the latent codes.

## 3 Model

### 3.1 Generative Adversarial Networks

The GAN[3] has two networks: a generator  $G$  that tries to generate real data given noise  $z \sim p_z(z)$ , and a discriminator  $D \in [0, 1]$  that classifies the real data  $x \sim p_{data}(x)$  and the fake data  $G(z)$ . The  $D(x)$  represents probability of  $x$  being a real data. The objective of  $G$  is to fit the true data distribution deceiving  $D$  by playing the following minimax game:

$$\min_{\theta_G} \max_{\theta_D} \mathbb{E}_{x \sim p_{data}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))], \quad (1)$$

where  $\theta_G$  and  $\theta_D$  are parameters of  $G$  and  $D$ , respectively. The parameters are updated by stochastic gradient descent (or ascent) algorithms.

### 3.2 Conditioned GAN

The objective of conditioned GANs is to fit conditional probability distribution  $p(x|c)$  where  $c$  is condition that describes  $x$ . The objective is to learn  $p(x|c)$  correctly from a labeled dataset  $(x_1, c_1), (x_2, c_2), \dots, (x_n, c_n)$ . The generator  $G(z, c)$  of conditioned GAN has additional input  $c$ , and all generators used in our experiment take  $[x, c]$  as an input where [...] means vector concatenation. We can add  $c$  into the input of  $D(x, c)$  as well, or put regularization terms to guide the generator. If we use  $D(x, c)$ , then the objective of the conditioned GAN is to optimize the equation 1 for each  $c$ :

$$\min_{\theta_G} \max_{\theta_D} \mathbb{E}_{c \sim p_{data}(c)} [\mathbb{E}_{x \sim p_{data}(x|c)} [\log D(x, c)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z, c), c))]], \quad (2)$$




Figure 1: Illustration of the four conditioned GANs: (a) CGAN, (b) FCGAN, (c) SBP, and (d) IRGAN. For simplicity of visualization, only a single pixel of  $x$  is visualized, virtually consisting of 5 channels.

where  $D(x, c)$  represents probability of  $x$  being a real data from condition  $c$ . We can also train  $G(x, c)$  to fit  $p(x|c)$  by putting a regularization term  $R(G)$  to the objective function:

$$\min_{\theta_G} \max_{\theta_D} \mathbb{E}_{x \sim p_{data}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z), c \sim p_{data}(c)} [\log(1 - D(G(z, c)))] + R(G). \quad (3)$$

### 3.3 Concatenation Methods

We compare our proposed models with two commonly used conditioned GANs. The first one is a conditional GAN[10] (CGAN) that concatenates  $c$  into  $x$  of  $D$ , and the second one is a fully conditional GAN (FCGAN) that concatenates  $c$  into every layer of  $D$  including  $x$ . When the dimension of a layer in  $D$  is  $n \times n \times d$  where  $d$  is the size of depth, we replicate the vector  $c$  spatially to match the size  $n \times n$  of the feature map and perform a depth concatenation. Both models, including ours, use the same structure for  $G(z, c)$  where  $c$  is only concatenated to  $z$ .

### 3.4 Spatial Bilinear Pooling

We propose Spatial Bilinear Pooling (SBP) that provides multiplicative interaction between all elements of two vectors. When the dimension of an image is  $n \times n \times d$ , the SBP performs cross product for each pixel ( $1 \times 1 \times d$ ) of the image with  $c$  and then gathers the resulting vectors spatially to make a new image.

### 3.5 Information Retrieving GAN

The Information Retrieving GAN (IRGAN) has an approximator  $Q(c|x)$  that measures  $p(c|x)$  from  $x$ . The generator  $G(z, c)$  of IRGAN maximizes  $Q(c|G(z, c))$ , fitting the conditional distribution  $p(x|c)$ . In our experiments, the approximators are pre-trained classifiers for MNIST and CIFAR-10. IRGAN optimizes the equation 3 where  $R(G)$  is defined by the lower bound[13] of the mutual information  $I(c; G(z, c))$ :

$$R(G) = -\lambda \mathbb{E}_{c \sim p_{data}(c), x \sim G(z, c)} [\log Q(c|x)]. \quad (4)$$

In this paper we assume that the entropy  $H(c)$  is constant and omit it from  $R(G)$  for simplicity. The only difference from [13] is that  $c$  is sampled from the data distribution rather than a pre-defined distribution, since  $c$  are explicitly given from the datasets.

## 4 Experiments

We trained conditioned GANs on MNIST and CIFAR-10, and utilized the techniques proposed by DCGAN[4]. All generators have identical structure, and the variations only exist in the discriminators and the auxiliary networks. We measured log-likelihood of the test set data for each  $c$  by fitting

Table 1: Parzen window-based log-likelihood estimates on MNIST.

Model	Label of MNIST									
	0	1	2	3	4	5	6	7	8	9
CGAN	124.7	358.0	35.7	56.9	106.7	-10.6	113.7	183.5	53.1	166.2
FCGAN	28.7	531.4	-11.9	52.8	101.1	51.7	81.6	221.0	20.0	187.1
SBP	153.0	643.7	<b>81.6</b>	124.6	173.2	121.4	197.2	287.5	83.1	248.7
IRGAN	<b>154.0</b>	<b>749.7</b>	68.7	<b>137.3</b>	<b>208.8</b>	<b>130.8</b>	<b>214.8</b>	<b>296.5</b>	<b>97.7</b>	<b>268.4</b>

Table 2: Parzen window-based log-likelihood estimates on CIFAR-10.

Model	Label of CIFAR-10									
	airplane	car	bird	cat	deer	dog	frog	horse	ship	truck
CGAN	684.0	417.4	969.5	603.8	1064.5	547.8	929.9	554.9	771.1	373.5
FCGAN	760.2	384.7	1018.9	566.3	1099.5	538.5	<b>985.9</b>	558.5	793.2	383.7
SBP	<b>847.7</b>	<b>471.3</b>	<b>1057.2</b>	<b>625.9</b>	<b>1180.9</b>	577.1	974.7	<b>606.8</b>	<b>872.9</b>	<b>460.6</b>
IRGAN	721.8	391.1	1038.0	561.1	1027.9	<b>586.4</b>	831.4	509.0	744.4	365.9

a Gaussian Parzen window to the samples generated from  $G(z, c)$ [3]. The best standard deviation  $\sigma$  of the Gaussian was chosen by the validation set.

The results are shown in Table 1, 2. The IRGAN showed the best result on MNIST, but it didn't performed well on CIFAR-10 due to inaccurate classifier  $Q(c|x)$  whose classification accuracy was 81%. We found that SBP showed stable results on both datasets and most of the labels, surpassing CGAN and FCGAN.

## 5 Conclusions and Future Works

In this paper we proposed effective models that can fit conditional distribution  $p(x|c)$  more accurately. However, we need more experiments on other complex conditions such as multi-labels, text descriptions, and style embedding to verify our models. Since cross product of two long vectors is inefficient, compressing bilinear pooling is a promising alternative to the spatial bilinear pooling when conditions are high dimensional vectors.

## References

- [1] Yoshua Bengio. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009.
- [2] L. Theis, A. van den Oord, and M. Bethge. A note on the evaluation of generative models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016.
- [3] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- [4] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [5] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *arXiv preprint arXiv:1606.03498*, 2016.
- [6] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*, 2016.

- [7] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *International Conference on Machine Learning (ICML)*, 2016.
- [8] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016.
- [9] Junbo Zhao, Michael Mathieu, and Yann LeCun. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*, 2016.
- [10] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [11] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 1558–1566, 2016.
- [12] Dong Huk Park, Daylen Yang, Akira Fukui, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *Conference on Empirical Methods in Natural Language Processing*, 2016.
- [13] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *arXiv preprint arXiv:1606.03657*, 2016.
- [14] Xiaolong Wang and Abhinav Gupta. Generative image modeling using style and structure adversarial networks. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV*, pages 318–335, 2016.
- [15] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 689–696, 2011.
- [16] Bolei Zhou, Yuandong Tian, Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. Simple baseline for visual question answering. *arXiv preprint arXiv:1512.02167*, 2015.
- [17] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. *arXiv preprint arXiv:1511.02274*, 2015.
- [18] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnn models for fine-grained visual recognition. In *International Conference on Computer Vision (ICCV)*, 2015.
- [19] Yang Gao, Oscar Beijbom, Ning Zhang, and Trevor Darrell. Compact bilinear pooling. In *CVPR*, 2016.
- [20] Ninh Pham and Rasmus Pagh. Fast and scalable polynomial kernels via explicit feature maps. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 239–247. ACM, 2013.
- [21] David Barber Felix Agakov. The im algorithm: a variational approach to information maximization. In *Advances in Neural Information Processing Systems 16: Proceedings of the 2003 Conference*, volume 16, page 201. MIT Press, 2004.