Multi-Player Generative Adversarial Networks

Samarth Gupta
Massachusetts Institute of Technology
samarthg@mit.edu

***Abstract*—In unsupervised learning, Generative Adversarial Networks (GANs) [1] have become one of the most popular techniques to implicitly learn the underlying distribution of a particular dataset. The problem is formulated as a two-player zero-sum game between a Generator network $\mathcal{G}$ and a Discriminator network $\mathcal{D}$. The Generator network aims to generate synthetic(fake) images from random noise and the Discriminator aims to correctly classify between synthetic and natural images. In this study we aim to explore the possibility of using generalized zero-sum games or multi-player zero-sum games for formulating and training GANs. Our results indicate that using a multi-player setup, the severity of mode collapse problem in GANs can be significantly reduced to generate a diverse set of synthetic images. We demonstrate this on different synthetic toy distributions and real-world datasets like MNIST & CIFAR10.**

***Index Terms*—Generative Adversarial Networks, Generalized zero-sum games, Unsupervised learning, Deep Learning**

I. Introduction

The famous min-max theorem by Neumann [2] for 2 player zero-sum games is one of the most important and celebrated results in game theory. It turns out that doing a similar analysis for multi-player games is quite difficult and the results of two player games do not translate easily to multi-player setting. However recently, Cai and Daskalakis [3] and Cai et al. [4] have proved some interesting results for multi-player zero-sum games. In this study, we aim to extend the 2 player GAN setup to a multi-payer setting following the the key results from [3] and [4].

A multi-player separable game can be viewed as a graph, where the nodes represent individual players and edges correspond to a two person game(not necessarily a zero-sum) being played between the nodes/players. Every player has a fixed set of strategies and has to choose from this set in all games it plays with its adjacent nodes/players. The payoff of a node/player is given by the sum of the payoffs of all games on its adjacent edges for a particular strategy profile. The game is *zero-sum* if for all strategy profiles, the sum of payoffs of all players add up to zero.

A multi-player game of the above form is definitely zero-sum if every game being played on an edge is zero-sum. But this is not a necessary condition, there may be a case when the individual games are not zero-sum but the overall game is still zero-sum. The main key findings in [2] are as follows:

1) LP formulation is possible to solve a multi-player zero-sum game.
2) An efficient LP formulation based algorithm to identify whether a generic multi-player game is zero-sum or not.
3) Nash equilibrium for such games can also be computed using no-regret learning algorithms like fictitious play, hedging etc.
4) Unlike two player games, the payoff of each player is not unique in each Nash equilibrium, but the sum of the payoff's of all players should add upto zero. Therefore, it can be viewed as a *closed-system of payoffs*.

Another interesting feature of these games is that the zero-sum nature of the game is a *global property* rather than a local property as individual games on the edges may be non-zero sum.

II. Multi player GANs

For simplicity, we limit ourselves to a star-network setting with $n+1$ nodes, where the central node acts as a Generator $\mathcal{G}$ and all the other nodes act as a Discriminator $\mathcal{D}_i \quad \forall i \in \{1,..,n\}$. We know that the underlying distribution which we are trying to learn is highly multi-modal. The key idea is that in each individual game between $\mathcal{G}$ & $\mathcal{D}_i$, we only use a subset of the original dataset. In this setting there can be two different cases: First where each Discriminator has images/(real data) from only one class and second, where the dataset is randomly distributed amongst different Discriminators.

Algorithm 1 Multi player GAN Training Algorithm

for no of epochs **do**
  for no_of_batches in each epoch **do**

    for k in no_of_Discriminators **do**
      Sample minibatch of m noise samples $z^1, ..., z^m$ from noise prior $p_g(z)$
      Sample minibatch of m examples $x^1, ..., x^m$ from **dataset** D_k
      Update the discriminator D_k by ascending its stochastic gradient:

$$\nabla_{\theta_{d_k}} \frac{1}{m} \sum_{i=1}^{m} [logD_k(x^i) + log(1 - D_k(G(z^i)))]$$

    end for
    Sample K minibatches of size m noise samples $\{z_1^1, ..., z_1^m\}.....\{z_K^1, ..., z_K^m\}$ from noise prior $p_g(z)$
    Update the generator by descending its stochastic gradient:

$$\nabla_{\theta_g} \frac{1}{k} \frac{1}{m} \sum_{k=1}^{K} \sum_{i=1}^{m} [log(1 - D_k(G(z_k^i)))]$$

  end for
end for

978-1-5386-5989-2/18/$31.00 ©2018 IEEE

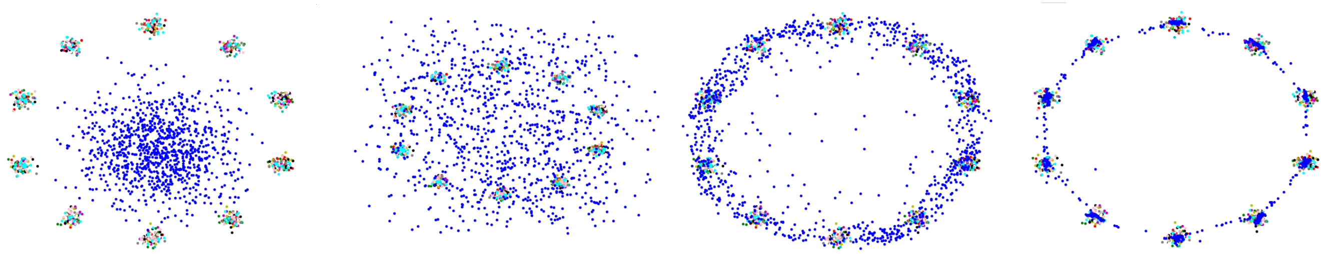

Fig. 1: Images generated during training on a mixture of 10 Gaussians

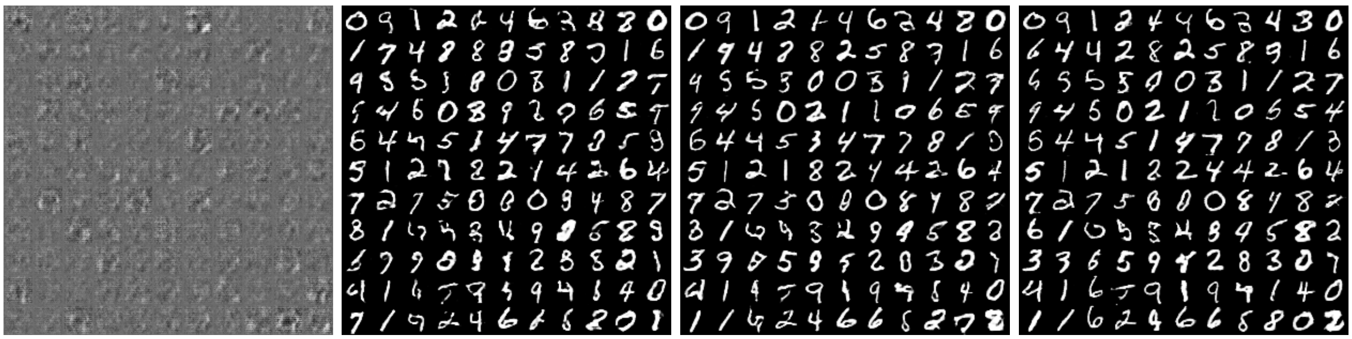

Fig. 2: Images generated during training on MNIST

Fig. 3: Images generated during training on CIFAR10

III. Results

We tested our proposed algorithm on a number of different toy and real-world datasets. We observed that for case 1 setting where each Discriminator has data points only from one particular class/mode suffers severely from mode collapse. However, case 2 where the dataset is randomly distributed amongst all Discriminators, seems to be quite robust to mode-collapse. Figure 1,2,3 shows the performance of case 2. We use a classification-based method recently proposed [5] to evaluate our performance on real-datasets. Figure 4 shows the probability of generating an image for a particular class. Our results are inline with MIX-DCGAN approach [6], where a mixture of Generators and Discriminators is used.

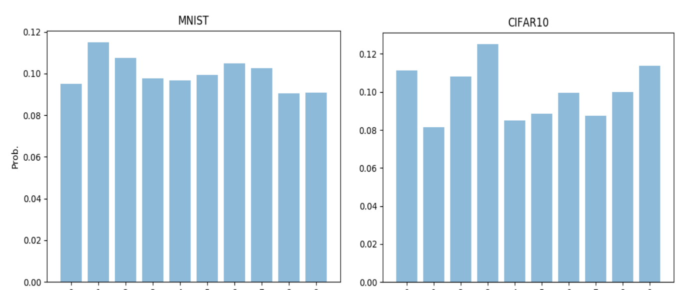

Fig. 4: Classification-based performance (Sample size=50k)

References

[1] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets
[2] John von Neumann. Zur Theorie der Gesellshaftsspiele. Mathematische Annalen 100:295-320. 1928.
[3] Yang Cai and Constantinos Daskalakis. On Minmax Theorems for Multiplayer Games.
[4] Yang Cai, Ozan Candogan, Constantinos Daskalakis and Christos Papadimitriou. A Multiplayer Generalization of the MinMax Theorem.
[5] Shibani Santurkar, Ludwig Schmidt, Aleksander Madry. A Classification-based perspective on GAN distributions.
[6] Sanjeev Arora, Rong Ge, Yingyu Liang,Tengyu Ma, Yi Zhang. Generalization and Equilibrium in Generative Adversarial Nets (GANs).