



CapsuleGAN: Generative Adversarial Capsule Network

Ayush Jaiswal^(✉), Wael AbdAlmageed, Yue Wu, and Premkumar Natarajan

USC Information Sciences Institute, Marina del Rey, CA, USA
{ajaiswal,wamageed,yue_wu,pnataraj}@isi.edu

Abstract. We present Generative Adversarial Capsule Network (CapsuleGAN), a framework that uses capsule networks (CapsNets) instead of the standard convolutional neural networks (CNNs) as discriminators within the generative adversarial network (GAN) setting, while modeling image data. We provide guidelines for designing CapsNet discriminators and the updated GAN objective function, which incorporates the CapsNet margin loss, for training CapsuleGAN models. We show that CapsuleGAN outperforms convolutional-GAN at modeling image data distribution on MNIST and CIFAR-10 datasets, evaluated on the generative adversarial metric and at semi-supervised image classification.

Keywords: Capsule networks · Generative adversarial networks

1 Introduction

Generative modeling of data is a challenging machine learning problem that has garnered tremendous interest recently, partly due to the invention of generative adversarial networks (GANs) [5] and its several sophisticated variants¹. A GAN model is typically composed of two neural networks; (1) a generator that attempts to transform samples drawn from a prior distribution to samples from a complex data distribution with much higher dimensionality, and (2) a discriminator that decides whether the given sample is real or from the generator's distribution. The two components are trained by playing an adversarial game. GANs have shown great promise in modeling highly complex distributions underlying real world data, especially images. However, they are notorious for being difficult to train and have problems with stability, vanishing gradients, mode collapse and inadequate mode coverage [4, 16, 19]. Consequently, there has been a large amount of work towards improving GANs by using better objective functions [1, 2, 6], sophisticated training strategies [19], using structural hyperparameters [15, 16] and adopting empirically successful tricks².

Radford et al. [16] provide a set of architectural guidelines, formulating a class of convolutional neural networks (CNNs) that have since been extensively

¹ <https://github.com/hindupuravinash/the-gan-zoo>.

² <https://github.com/soumith/ganhacks>.

used to create GANs (referred to as Deep Convolutional GANs or DCGANs) for modeling image data and other related applications [10, 17]. More recently, however, Sabour et al. [18] introduced capsule networks (CapsNets) as a powerful alternative to CNNs, which learn a more *equivariant* representation of images that is more robust to changes in pose and spatial relationships of parts of objects in images [7] (information that CNNs lose during training, by design). Inspired by the working mechanism of optic neurons in the human visual system, capsules were first introduced by Hinton et al. [7] as locally invariant groups of neurons that learn to recognize visual entities and output activation vectors that represent both the presence of those entities and their properties relevant to the visual task (such as object classification). The training algorithm of CapsNets involves a routing mechanism between capsules in successive layers of the network that imitates hierarchical communication of information across neurons in human brains that are responsible for visual perception and understanding.

The initial intuition behind the design of deep neural networks was to imitate human brains for modeling hierarchical recognition of features, starting from low-level attributes and progressing towards complex entities. CapsNets capture this intuition more effectively than CNNs because they have the aforementioned in-built explicit mechanism that models it. CapsNets have been shown to outperform CNNs on MNIST digit classification and segmentation of overlapping digits [18]. This motivates the question whether GANs can be designed using CapsNets (instead of CNNs) to improve their performance.

We propose Generative Adversarial Capsule Network (CapsuleGAN) as a framework that incorporates capsules within the GAN framework. In particular, CapsNets are used as discriminators in our framework as opposed to the conventionally used CNNs. We show that CapsuleGANs perform better than CNN-based GANs at modeling the underlying distribution of MNIST [13] and CIFAR-10 [12] datasets both qualitatively and quantitatively using the generative adversarial metric (GAM) [9] and at semi-supervised classification using unlabeled GAN-generated images with a small number of labeled real images.

The rest of the paper is organized as follows. Section 2 discusses related work. In Sect. 3 we provide a brief introduction to GANs and CapsNets. Section 4 describes our CapsuleGAN framework along with implementation guidelines. Qualitative and quantitative analyses of our model are presented in Sect. 5. Section 6 concludes the paper and provides directions for future research.

2 Related Work

GANs were originally implemented as feedforward multi-layer perceptrons, which did not perform well on generating complex images like those in the CIFAR-10 dataset [12]. They suffered from mode collapse and were highly unstable to train [16, 19]. In an attempt to solve these problems, Radford et al. [16] presented a set of guidelines to design GANs as a class of CNNs, giving rise to DCGANs, which have since been a dominant approach to GAN network architecture design. Im et al. [8] later proposed the use of Recurrent Neural Networks instead of CNNs

as generators for GANs, creating a new class of GANs referred to as Generative Recurrent Adversarial Networks or GRANs. On a related note, Odena et al. [15] proposed an architectural change to GANs in the form of a discriminator that also acts as a classifier for class-conditional image generation. This approach for designing discriminators has been a popular choice for conditional GANs [14] recently. Our work is similar in line with [15] in the sense that we propose an architectural change to discriminators. We propose to transition from designing GAN discriminators as CNNs to formulating them as CapsNets, creating a new class of GANs called CapsuleGANs. This idea can be extended to encoder-based GANs like BiGAN [3] where the encoder can be modeled as a CapsNet also.

3 Preliminaries

3.1 Generative Adversarial Networks

Goodfellow et al. [5] introduced GANs as a framework for generative modeling of data through learning a transformation from points belonging to a simple prior distribution ($\mathbf{z} \sim p_z$) to those from the data distribution ($\mathbf{x} \sim p_{data}$). The framework is composed of two models that play an adversarial game: a generator and a discriminator. While the generator attempts to learn the aforementioned transformation $G(\mathbf{z})$, the discriminator acts as a critic $D(\cdot)$ determining whether the sample provided to it is from the generator’s output distribution ($G(\mathbf{z}) \sim p_G$) or from the data distribution ($\mathbf{x} \sim p_{data}$), thus giving a scalar output ($y \in \{0, 1\}$). The goal of the generator is to fool the discriminator by generating samples that resemble those from the real data while that of the discriminator is to accurately distinguish between real and generated data. The two models, typically designed as neural networks, play an adversarial game with the objective as shown in Eq. 1.

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \quad (1)$$

3.2 Capsule Networks

The concept of capsules was first introduced by Hinton et al. [7] as a method for learning robust unsupervised representation of images. Capsules are locally invariant groups of neurons that learn to recognize the presence of visual entities and encode their properties into vector outputs, with the vector length (limited to being between zero and one) representing the presence of the entity. For example, each capsule can learn to identify certain objects or object-parts in images. Within the framework of neural networks, several capsules can be grouped together to form a capsule-layer where each unit produces a vector output instead of a (conventional) scalar activation.

Sabour et al. [18] introduced a routing-by-agreement mechanism for the interaction of capsules within deep neural networks with several capsule-layers, which works by pairwise determination of the passage of information between capsules in successive layers. For each capsule $h_i^{(l)}$ in layer l and each capsule $h_j^{(l+1)}$ in

the layer above, a coupling coefficient c_{ij} is adjusted iteratively based on the agreement (cosine similarity) between h_i 's prediction of the output of h_j and its actual output given the product of c_{ij} and h_i 's activation. Thus, the coupling coefficients inherently decide how information flows between pairs of capsules. For a classification task involving K classes, the final layer of the CapsNet can be designed to have K capsules, each representing one class. Since the length of a capsule's vector output represents the presence of a visual entity, the length of each capsule in the final layer ($\|\mathbf{v}_k\|$) can then be viewed as the probability of the image belonging to a particular class (k). The authors introduce a margin loss L_M for training CapsNets for multi-class classification, as show in Eq. 2:

$$L_M = \sum_{k=1}^K T_k \max(0, m^+ - \|\mathbf{v}_k\|)^2 + \lambda(1 - T_k) \max(0, \|\mathbf{v}_k\| - m^-)^2 \quad (2)$$

where T_k represents target labels, $m^+ = 0.9$, $m^- = 0.1$ and $\lambda = 0.5$, a down-weighting factor for preventing initial learning from shrinking the lengths of the capsule outputs in the final layer. The authors also add regularization to the network in the form of a weighted image reconstruction loss, where the vector outputs \mathbf{v}_k of the final layer are presented as inputs to the reconstruction network.

4 Generative Adversarial Capsule Networks

GANs have been mostly used for modeling the distribution of image data and associated attributes, as well as for other image-based applications like image-to-image translation [10] and image synthesis from textual descriptions [17]. The generator and the discriminator have conventionally been modeled as deep CNNs following the DCGAN guidelines [16]. We follow this convention in designing the CapsuleGAN generator as a deep CNN. However, motivated by the stronger intuition behind and the superior performance of CapsNets with respect to CNNs [18], we design the proposed CapsuleGAN framework to incorporate capsule-layers instead of convolutional layers in the GAN discriminator, which fundamentally performs a two-class classification task.

The CapsuleGAN discriminator is similar in architecture to the CapsNet model presented in [18]. CapsNets, in general, have a large number of parameters because, firstly, each capsule produces a vector output instead of a single scalar and, secondly, each capsule has *additional* parameters associated with all the capsules in the layer above it that are used for making predictions about their outputs. However, it is necessary to keep the number of parameters in the CapsuleGAN discriminator low due to two reasons: (1) CapsNets are very powerful models and can easily start harshly penalizing the generator early on in the training process, which will cause the generator to either fail completely or suffer from mode collapse, and (2) current implementations of the dynamic routing algorithm are slow to run. It is important to note that first reason for

keeping the number of parameters of the CapsNet low falls in line with the popular design of convolutional discriminators as relatively shallow neural networks with low numbers of relatively large-sized filters in their convolutional layers.

The final layer of the CapsuleGAN discriminator contains a single capsule, the length of which represents the probability whether the discriminator’s input is a real or a generated image. We use margin loss L_M instead of the conventional binary cross-entropy loss for training our CapsuleGAN model because L_M works better for training CapsNets. Therefore, the objective of CapsuleGAN can be formulated as shown in Eq. 3.

$$\begin{aligned} & \min_G \max_D V(D, G) \\ & = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [-L_M(D(\mathbf{x}), \mathbf{T} = \mathbf{1})] + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [-L_M(D(G(\mathbf{z})), \mathbf{T} = \mathbf{0})] \quad (3) \end{aligned}$$

In practice, we train the generator to minimize $L_M(D(G(\mathbf{z})), \mathbf{T} = \mathbf{1})$ instead of minimizing $-L_M(D(G(\mathbf{z})), \mathbf{T} = \mathbf{0})$. This essentially eliminates the down-weighting factor λ in L_M when training the generator, which does not contain any capsules.

5 Experimental Evaluation

We evaluate the performance of CapsuleGANs at randomly generating images through a series of experiments as described below, in which we compare CapsuleGANs with convolutional GANs both qualitatively and quantitatively. We implement both the GAN models with the same architecture for their generators. Both convolutional GAN and the proposed CapsuleGAN models are implemented using the publicly available `keras-adversarial`³ and `CapsNet-Keras`⁴ packages.

5.1 Data

We provide results of our experiments on MNIST and CIFAR-10 datasets. The MNIST dataset consists of 28×28 sized grayscale images of handwritten digits. The CIFAR-10 dataset contains 32×32 color images grouped into ten classes: airplane, automobile, bird, cat, deer, dog, frog, horse, ship and truck.

5.2 Visual Quality of Randomly Generated Images

We qualitatively compare images generated randomly using both GAN and CapsuleGAN. Figures 1a and b show images generated using the standard convolutional-GAN and CapsuleGAN, respectively, on the MNIST dataset. Qualitatively, both CapsuleGAN and the standard convolutional-GAN produce crisp images of similar quality, that sometimes do not resemble any digit. However, the image-grid generated using GAN seems to have less diversity in terms

³ <https://github.com/bstriner/keras-adversarial>.

⁴ <https://github.com/XifengGuo/CapsNet-Keras>.

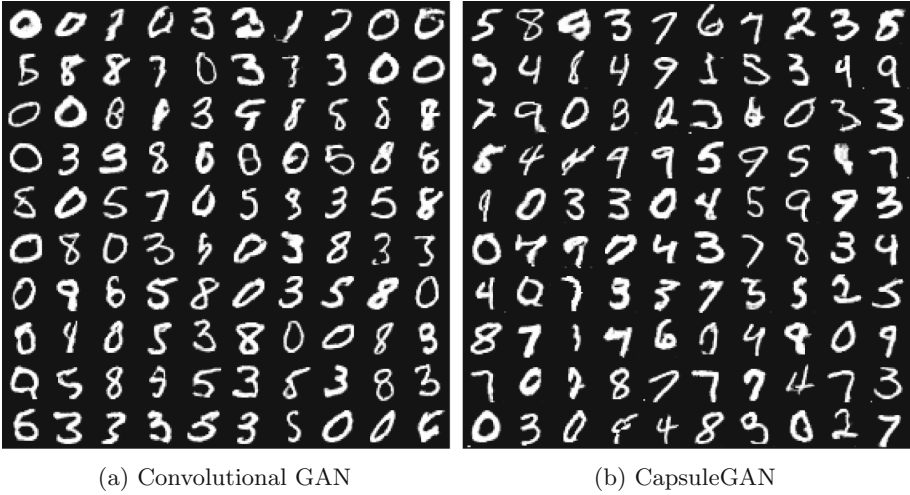


Fig. 1. Randomly generated MNIST images

of generated classes of digits. Figures 2a and b show the results of this experiment on the CIFAR-10 dataset. Both the models produce diverse sets of images but images generated using CapsuleGAN look cleaner and crisper than those generated using convolutional-GAN. We provide results of our quantitative evaluation in the following subsections for deeper analyses of the image generation performance.

5.3 Generative Adversarial Metric

Im et al. [9] introduced the generative adversarial metric (GAM) as a pairwise comparison metric between GAN models by pitting each generator against the opponent’s discriminator, i.e., given two GAN models $M_1 = (G_1, D_1)$ and $M_2 = (G_2, D_2)$, G_1 engages in a battle against D_2 while G_2 against D_1 . The ratios of their classification errors on real test dataset and on generated samples are then calculated as r_{test} and $r_{samples}$. Following their implementation⁵, in practice, the ratios of classification accuracies are calculated instead of errors to avoid numerical problems, as shown in Eqs. 4 and 5

$$r_{samples} = \frac{A(D_{GAN}(G_{CapsuleGAN}(\mathbf{z})))}{A(D_{CapsuleGAN}(G_{GAN}(\mathbf{z})))} \quad (4)$$

$$r_{test} = \frac{A(D_{GAN}(\mathbf{x}_{test}))}{A(D_{CapsuleGAN}(\mathbf{x}_{test}))} \quad (5)$$

Therefore, for CapsuleGAN to win against GAN, both $r_{samples} < 1$ and $r_{test} \simeq 1$ must be satisfied. In our experiments, we achieve $r_{samples} = 0.79$

⁵ <https://github.com/jiwoongim/GRAN/battle.py>.

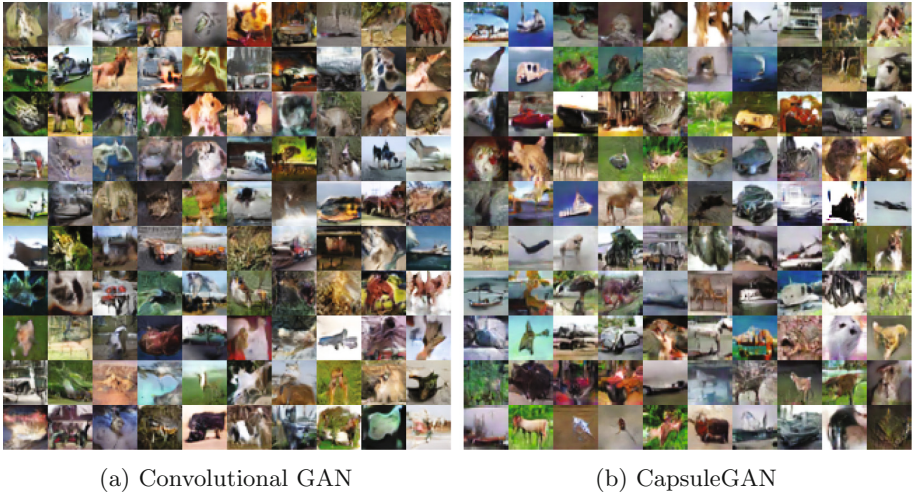


Fig. 2. Randomly generated CIFAR-10 images

Table 1. Results of semi-supervised classification - MNIST

Model	Error rate		
	n = 100	n = 1,000	n = 10,000
Convolutional GAN	0.2900	0.1539	0.0702
CapsuleGAN	0.2724	0.1142	0.0531

and $r_{test} = 1$ on the MNIST dataset and $r_{samples} = 1.0$ and $r_{test} = 0.72$ on the CIFAR-10 dataset. Thus, on this metric, CapsuleGAN performs better than convolutional GAN on the MNIST dataset but the two models tie on the CIFAR-10 dataset.

5.4 Semi-supervised Classification

We evaluate the performance of the convolutional GAN and the proposed CapsuleGAN on semi-supervised classification. In this experiment, we randomly generate 50,000 images using both GAN and CapsuleGAN. We use the Label Spreading algorithm [20] with the generated images as the unlabeled examples and n real labeled samples, with $n \in \{100, 1000, 10000\}$. We use the `scikit-learn`⁶ package for these experiments. Table 1 shows the results of our experiments on MNIST while Table 2 shows those on CIFAR-10. The error rates are high in most experimental settings because we provide raw pixel values as features to the classification algorithm. However, this allows us to more objectively compare the two models without being biased by feature extraction methods.

⁶ <http://scikit-learn.org/>.

Table 2. Results of semi-supervised classification - CIFAR-10

Model	Error rate		
	n = 100	n = 1,000	n = 10,000
Convolutional GAN	0.8305	0.7587	0.7209
CapsuleGAN	0.7983	0.7496	0.7102

The results show that the proposed CapsuleGAN consistently outperforms convolutional GAN for all the tested values of n with a margin of 1.7 – 3.97 % points for MNIST and 0.91 – 3.22 % points for CIFAR-10. Thus, CapsuleGAN generates images that are more similar to real images and more diverse than those generated using convolutional GAN, leading to better semi-supervised classification performance on the test dataset.

6 Discussion and Future Work

Generative adversarial networks are extremely powerful tools for generative modeling of complex data distributions. Research is being actively conducted towards further improving them as well as making their training easier and more stable. Motivated by the success of CapsNets over CNNs at image-based inference tasks, we presented the generative adversarial capsule network (CapsuleGAN), a GAN variant that incorporates CapsNets instead of CNNs as discriminators when modeling image data. We presented guidelines for designing CapsuleGANs as well as an updated objective function for training CapsuleGANs. We showed that CapsuleGANs outperform convolutional-GANs on the generative adversarial metric and at semi-supervised classification with a large number of unlabeled generated images and a small number of real labeled ones, on MNIST and CIFAR-10 datasets. This indicates that CapsNets should be considered as potential alternatives to CNNs for designing discriminators and other inference modules in future GAN models.

We plan to conduct theoretical analysis of the use of margin loss within the GAN objective. We purposefully did not incorporate many GAN training tricks to fairly evaluate our contributions. The results presented in this paper motivate the use of CapsNets as opposed to CNNs for encoders in GAN variants like BiCoGAN [11]. We see this as an important direction for future research.

Acknowledgements. This work is based on research sponsored by the Defense Advanced Research Projects Agency under agreement number FA8750-16-2-0204. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein GAN. arXiv preprint [arXiv:1701.07875](https://arxiv.org/abs/1701.07875) (2017)
2. Berthelot, D., Schumm, T., Metz, L.: Began: boundary equilibrium generative adversarial networks. arXiv preprint [arXiv:1703.10717](https://arxiv.org/abs/1703.10717) (2017)
3. Donahue, J., Krähenbühl, P., Darrell, T.: Adversarial feature learning. In: International Conference on Learning Representations (2017)
4. Durugkar, I., Gemp, I., Mahadevan, S.: Generative multi-adversarial networks. In: International Conference on Learning Representations (2017)
5. Goodfellow, I., et al.: Generative adversarial nets. In: Advances in Neural Information Processing Systems, pp. 2672–2680 (2014)
6. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of Wasserstein GANs. In: Advances in Neural Information Processing Systems, pp. 5769–5779 (2017)
7. Hinton, G.E., Krizhevsky, A., Wang, S.D.: Transforming auto-encoders. In: Honkela, T., Duch, W., Girolami, M., Kaski, S. (eds.) ICANN 2011. LNCS, vol. 6791, pp. 44–51. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21735-7_6
8. Im, D.J., Kim, C.D., Jiang, H., Memisevic, R.: Generating images with recurrent adversarial networks. arXiv preprint [arXiv:1602.05110](https://arxiv.org/abs/1602.05110) (2016)
9. Im, D.J., Kim, C.D., Jiang, H., Memisevic, R.: Generative adversarial metric (2016)
10. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017
11. Jaiswal, A., AbdAlmageed, W., Wu, Y., Natarajan, P.: Bidirectional Conditional Generative Adversarial Networks. arXiv preprint [arXiv:1711.07461](https://arxiv.org/abs/1711.07461) (2017)
12. Krizhevsky, A.: Learning multiple layers of features from tiny images (2009)
13. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**(11), 2278–2324 (1998)
14. Mirza, M., Osindero, S.: Conditional Generative Adversarial Nets. arXiv preprint [arXiv:1411.1784](https://arxiv.org/abs/1411.1784) (2014)
15. Odena, A., Olah, C., Shlens, J.: Conditional image synthesis with auxiliary classifier GANs. In: Precup, D., Teh, Y.W. (eds.) Proceedings of the 34th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol. 70, pp. 2642–2651. PMLR, International Convention Centre, Sydney, Australia, 06–11 August 2017. <http://proceedings.mlr.press/v70/odena17a.html>
16. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. In: International Conference on Learning Representations (2016)
17. Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text-to-image synthesis. In: Proceedings of The 33rd International Conference on Machine Learning (2016)
18. Sabour, S., Frosst, N., Hinton, G.E.: Dynamic routing between capsules. In: Advances in Neural Information Processing Systems, pp. 3859–3869 (2017)

19. Salimans, T., et al.: Improved techniques for training GANs. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 29, pp. 2234–2242. Curran Associates, Inc. (2016). <http://papers.nips.cc/paper/6125-improved-techniques-for-training-gans.pdf>
20. Zhou, D., Bousquet, O., Lal, T.N., Weston, J., Schölkopf, B.: Learning with local and global consistency. In: *Advances in Neural Information Processing Systems*, vol. 16. pp. 321–328. MIT Press (2004)