Self-Supervised Image-to-Text and Text-to-Image Synthesis

Anindya Sundar Das(✉) and Sriparna Saha

Department of Computer Science and Engineering
Indian Institute of Technology Patna, India
`anindyasd34, sriparna.saha@gmail.com`

Abstract. A comprehensive understanding of vision and language and their interrelation are crucial to realize the underlying similarities and differences between these modalities and to learn more generalized, meaningful representations. In recent years, most of the works related to Text-to-Image synthesis and Image-to-Text generation, focused on supervised generative deep architectures to solve the problems, where very little interest was placed on learning the similarities between the embedding spaces across modalities. In this paper, we propose a novel self-supervised deep learning based approach towards learning the cross-modal embedding spaces; for both image to text and text to image generations. In our approach, we first obtain dense vector representations of images using StackGAN-based autoencoder model and also dense vector representations on sentence-level utilizing LSTM based text-autoencoder; then we study the mapping from embedding space of one modality to embedding space of the other modality utilizing GAN and maximum mean discrepancy based generative networks. We, also demonstrate that our model learns to generate textual description from image data as well as images from textual data both qualitatively and quantitatively.

Keywords: Cross-modal · Semantic Space · Embedding Space · Maximum Mean Discrepancy · Mapping Networks.

1 Introduction

The web contains a multitude of images; most of the content images are unannotated. Describing the content of an image automatically using proper natural languages is a vital task, relevant to the area of both Natural Language Processing and Computer Vision, the impact of which could be significant, for example, it will help visually impaired people to have a better understanding of the content of images on the web using existing text-to-speech systems. It has many other important applications such as semantic visual search [6], or visual intelligence in chatbots [4]. The reverse problem is the generation of realistic images from human-written descriptions. Although notable progress has been made in generating visually realistic images, those are still far from this goal. The Generative Adversarial Network (GAN) [7] based models showed promising

results by generating many probable visual representations of a given textual description [22], [23], [31].

However, one major problem in both text to image synthesis and image to text generation is that recent state-of-the-art deep models are supervised learning-based that require annotated data. Most of the available web data is unlabeled and requires expensive human annotations. Similar problems have been addressed in machine translation where unsupervised machine translation [13] attains astounding results even in cases of low-resource languages and also performs reasonably well in the case of distant languages [14]. Motivated by the recent success in machine translation, we aim to investigate whether it is possible to learn the cross-modal embedding spaces between text and visual data in a self-supervised fashion.

*The major contributions of this paper can be summarized as follows: **i.** To the best of our knowledge, the current work is the first attempt in developing a self-supervised or unsupervised way of generating images from texts and texts from images. **ii.** Firstly a generative deep image autoencoder setup is developed for generating compelling image vector semantic space as per the reconstructions are concerned. Secondly, a recurrent neural network based autoencoder is developed which embeds sentences in textual semantic space. Finally, the mapping between cross-modal semantic spaces is established using both GANs and Maximum Mean Discrepancy (MMD) [32] based generative networks [16] which utilize the adversarial kernel learning technique. **iii.** Results on Caltech-UCSD Birds-200-2011 and Oxford-102 datasets illustrate that self-supervised Image-to-Text and Text-to-Image generation techniques can generate one modality given the other modality.*

2   Related Works

Caption generation using neural networks was first proposed in [12] which used a multimodal log-bi-linear model. In the paper [27], the authors used deep convolution neural networks (CNN) as an image encoder while RNN as a decoder that generates captions. The authors of [29], introduced an attention-based approach for image caption generation which uses convolutional networks for image feature extraction and attention-based RNN as decoder for caption generation. A deep visual semantic captioning model is proposed in [26] that takes the advantage of external sources and exploits the semantic information to generate captions and can describe the objects not present in image-caption datasets. A caption generation model that is based on the dependencies between caption words, image regions, and RNN language model has been proposed in [20]. The authors [15] showed in their work, that the meaningful style representation which is not well-described in text, but present in the images can be learned in an unsupervised manner. There have been several works on the generation of text for a given image, but there are also the ones that generate an image for a given text. Generative Adversarial Networks(GANs) [7] are proven to be useful in generating photo-realistic images [22] [30]. For text to image task, StackGAN-v2 [31] employs

a tree-like structure comprising of multiple generators and discriminators. It generates images at multiple scales at different branches of the tree.

Our work is as closely related to machine translation as it is related to image-caption generation. Recent work in unsupervised machine translation [13] uses a shared encoder-decoder architecture and maps the sentences from two different languages into common latent space. A work, related to unsupervised learning [16] that employs maximum mean discrepancy (MMD) [32] based on two-sample tests, minimizes the MMD distance between the two distributions to obtain the mapping. Building on the ideas and advances in these previous related works, we propose a novel architecture for cross-modal generations that utilizes GAN and MMD GANs [16] for unsupervised learning.

3   Methods

Our proposed self-supervised framework comprises of three sub-modules: 1) A deep StackGAN [31] based image autoencoder 2) LSTM-based Sequence-to-Sequence text autoencoder model 3) A cross-modal embedding space mapper that maps embedding from one semantic space to the other. We make our code publicly available[1].

3.1   StackGAN based Image Autoencoder

Our proposed model for image autoencoder is StackGAN based. StackGAN-v2 [31] takes a text embedding vector as input and produces images of increasing resolution at different branches of the network. We modified the architecture to make it an encoder-decoder based network as shown in Figure 1, where an encoder takes an image, extracts different features at different layers of the deep network to obtain an image embedding; this image embedding subsequently is fed as input to the original conditional StackGAN decoder model which reconstructs the image at the output of the conditional StackGAN thus working as an image autoencoder. For the encoder part of the autoencoder, we have used the pre-trained ResNet-50 [9]. We have redefined the last layer of ResNet-50 to obtain 1024 dimensional image embedding.

On the decoder side, the generated image is conditioned on the image-embedding vector. As described in Figure 1, firstly the ResNet encoder encodes image i into image embedding ψ_i. Then Conditional Augmentation technique [31] is applied on image embedding ψ_i to produce continuous latent space, yielding condition variable $\hat{c} \sim \mathcal{N}(\mu(\psi_i), \sigma(\psi_i))$. The following Kullback-Leibler (KL) divergence loss term is optimized during generator training which acts as regularization that ensures the smoothness of the conditioning variable distribution:

$$D_{KL}(\mathcal{N}(\mu(\psi_i), \sigma(\psi_i)) \| \mathcal{N}(0, 1)) \qquad (1)$$

For each generator, G_i, at different branches of the tree, the hidden feature, h_i, at i-th branch is computed as $h_0 = Fn_0(c, z)$ and $h_i = Fn_i(h_{i-1}, c)$ where Fn_i is
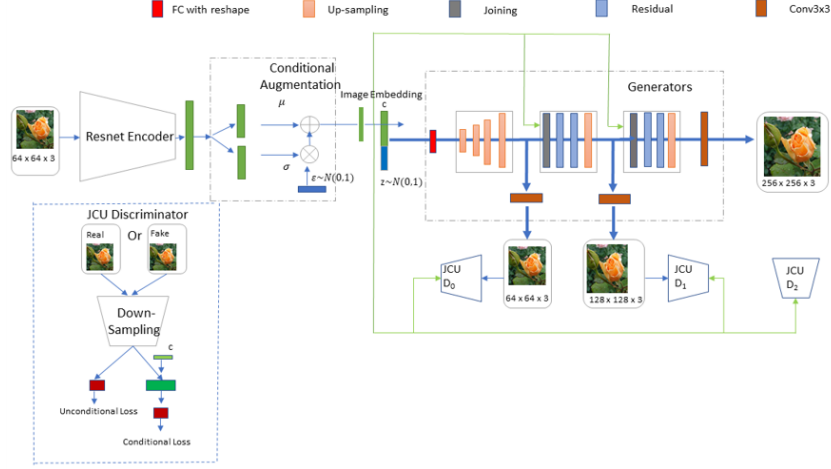
---

[1] https://github.com/anindyasdas/SelfSupervisedImageText

Fig. 1: Image autoencoder architecture: StackGAN Image Autoencoder. c is conditioning variable

neural network, $i = 0, 1, ..., n - 1$; n is the total number of branches. Generators at different stages produce images $u_i = G_i(hi)$ from low-to-high resolutions gradually adding more details. Both conditional loss and unconditional loss are being optimized while training the discriminator, D_i:

$$Loss_{D_i} = -\mathbb{E}_{x_i \sim p_{data_i}}[\log(D_i(x_i)] - \mathbb{E}_{u_i \sim p_{G_i}}[\log(1 - D_i(u_i)]$$
$$- \mathbb{E}_{x_i \sim p_{data_i}}[\log(D_i(x_i, c)] - \mathbb{E}_{u_i \sim p_{G_i}}[\log(1 - D_i(u_i, c)] \quad (2)$$

The unconditional loss dictates whether the image at discriminator input is fake or real; the conditional loss decides whether the generated image corresponds to the respective conditioning variable (Figure 1). During the training of the generators, the following loss function is being optimized:

$$Loss_{G_i} = -\mathbb{E}_{u_i \sim p_{G_i}}[\log(D_i(u_i)] - \mathbb{E}_{u_i \sim p_{G_i}}[\log(D_i(u_i, c)] \quad (3)$$

The final loss function for training the generators is given by

$$Loss_G = \sum_0^{n-1} Loss_{G_i} \quad (4)$$

3.2 LSTM-based Sequence-to-Sequence Text Autoencoder Model

We have used a single-layer LSTM encoder and decoder in our text autoencoder as a sequence-to-sequence model. The encoder is bidirectional LSTM with hidden

dimension 50, while the decoder is unidirectional with hidden size 100. The text encoder takes the word embedding vectors at each time step as inputs and generates corresponding hidden vectors. In this paper, we have considered all the hidden vectors to obtain the sentence embedding by max-pooling over all the hidden vectors. This latent vector is used to initialize the decoder LSTM network which regenerates the text description at the output.

3.3   Cross-modal Embedding Space Mapping Networks

Cross-modal embedding space mapping networks map from one modality embedding space to the other modality embedding space. The images and texts need not be paired, as both the networks minimize the distance between the two semantic distributions. We have employed two different architectures: One is GAN-based, and the other utilizes MMD-based generative networks.

GAN-based Cross-modal Embedding Space Mapping Networks
  In this architecture, simple GAN models are used, for both image-to-text and text-to-image conversions. The generator translates one modality embedding into the other modality embedding, and the discriminator determines whether two embedding distributions match or not.

MMD GAN-based Cross-modal Embedding Space Mapping Networks
  Maximum Mean Discrepancy (MMD) is a distance measure on the embedding probability space in Reproducing Kernel Hilbert Space (RKHS) [2]. Given two probability distributions P and Q and a continuous positive definite real-valued kernel k ($\mathcal{H}$ to be corresponding RKHS); the corresponding kernel means be defined as $\mu_p = \int k(.,x)dP(x)$ and $\mu_q = \int k(.,y)dQ(y)$, the distance $MMD(P,Q) = \|\mu_p - \mu_q\|$ measures the similarity between the two distributions, is known as MMD [8]. In this paper, we propose a mapping network based on MMD-GAN [16] that trains the generator g_θ to minimize MMD distance between two distributions, i.e. $\min_\theta M_k(\mathbb{P}_\mathcal{X}, \mathbb{P}_\theta)$, hence passes the hypothesis test.
  Here real data $x \sim \mathbb{P}_\mathcal{X}$, generator distribution $g_\theta(z) \sim \mathbb{P}_\theta$ and $\mathbb{P}_\mathcal{Z}$ is the base distribution such that $z \sim \mathbb{P}_\mathcal{Z}$ and M_k is the square of MMD distances :

$$
\begin{aligned}
M_k(\mathbb{P}_\mathcal{X}, \mathbb{P}_\theta) &= \|\mu_{\mathbb{P}_\mathcal{X}} - \mu_{\mathbb{P}_\theta}\|^2 \\
&= \mathbb{E}_{\mathbb{P}_\mathcal{X}}\{k(x,x')\} + \mathbb{E}_{\mathbb{P}_\theta}\{k(g_\theta(z), g_\theta(z'))\} \\
&\qquad - 2\mathbb{E}_{\mathbb{P}_\mathcal{X}, \mathbb{P}_\theta}\{k(x, g_\theta(z))\} \quad (5)
\end{aligned}
$$

4   Experiment

We have compared our MMD-based Cross-modal mapping network with a GAN-based mapping network, in which the generator generates a mapping from
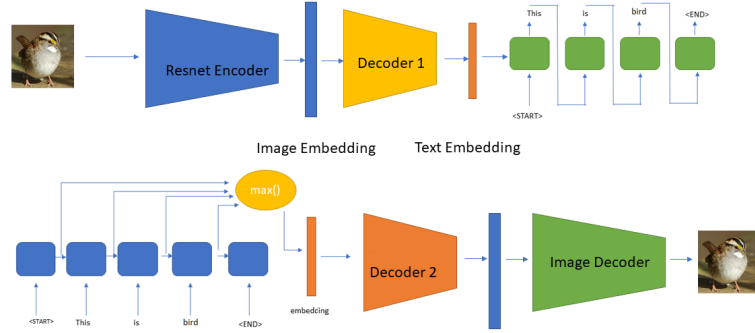
Fig. 2: End-to-end Networks: Self-supervised Image-to-Text Generator and Self-supervised Text-to-Image Generator

one embedding space to another, and the discriminator is a simple network to distinguish real and fake embeddings.

We first trained our image autoencoder (Section 3.1) for 600 epochs. We used pre-trained ResNet-50 encoder [9], the last layer of ResNet is re-defined as a fully connected layer with output dimension 1024, which is also our image vector dimension. During the training of the entire autoencoder, the weight of the ResNet, except the last layer, is kept fixed, while the parameters of the remaining network, are set as trainable. As the ResNet model is pre-trained on Imagenet dataset [5], we resized images to 224×224 and normalized images of our dataset using the mean and standard deviation of the Imagenet dataset.

Next, we pre-train our LSTM-based text autoencoder (Section 3.2) on 1 million sentences extracted from One Billion Word Benchmark dataset [3]. For initialization of our text autoencoder model, we have used top 50k 100-dimensional Glove embeddings [21] along with Glove embeddings of other unique words in caption datasets. The model was pre-trained for 50 epochs; this pre-trained model worked as an initializer; the model was then trained on all the captions of train sets. We have used Cross-Entropy between the generated and target sentences as the loss function.

After we trained both our autoencoders, the GAN-based Mapping Networks (Section 3.3) and the MMD-GAN Mapping Networks (Section 3.3) are trained separately using the unpaired image and text data, for both the image-to-text and the text-to-image conversions (Figure 2). The weights of the trained autoencoders are kept fixed while we train the mapping networks. In image-to-text setup, first, the image encoder encodes an input image into image embedding, which is then used as input to image-to-text embedding mapping networks, that generate corresponding text embedding, the batch of generated text embeddings are used as fake text embeddings while the text embeddings obtained directly from text encoder act as true text embeddings; fake embeddings and true embeddings are then used to train the generative networks. Likewise, the text-to-image setup is trained.

Baselines: We have compared our StackGAN image autoencoder with the following baselines:

ResNet autoencoder: In this architecture we fine tuned ResNet-50 as encoder and the decoder comprises of several layers of upsampling followed by deconvolution layer as used in [30]. The latent dimension is set to 1024.

Com-Rec autoencoder: This model is based on [11] which is completely CNN based. We added linear layers in the bottleneck to obtain dense vector (dimension 1024) latent space.

Compressive autoencoder: This architecture is based on model as discussed in the paper [25] which comprises of an encoder, a decoder and a probabilistic model. We have added a linear layer in the bottleneck to obtain dense 1024-dimensional vector embeddings.

For Mapping Networks, as this is a novel self-supervised approach with no existing baselines available for comparisons, we define the following baselines to evaluate the performance:

GAN-based Cross-modal Mapping Network: In this architecture simple GAN models are used for both the image-to-text and text-to-image conversions. The generator translates one modality embedding into the other modality embedding and the discriminator determines whether two embedding distributions match or not.

MMD GAN-based Cross-modal Mapping Network: This is statistics based network that uses MMD [8] as objective function (Section 3.3).

Dataset: For the task, we have used publicly available standard Caltech-UCSD Birds-200-2011 Dataset [28] and Oxford-102 Flower Dataset. The CUB-200-2011 contains a total of 11,788 images of 200 bird species, with multiple annotations per image. All attributes are visual; mostly related to the shape, color, or pattern of the particular part. We split the dataset into train set (8,855 images, 88540 captions, 150 classes) and test set (2,933 images, 29330 captions, 50 classes) such that the respective classes do not overlap. The object-image size ratio in the dataset is less than 0.5 for 80% of the images, so we pre-processed the images to maintain the size ratio at values greater than 0.75.

The Oxford-102 dataset contains a total of 8,189 images of 102 categories of flowers. The dataset is split into train set (7,034 images, 70340 captions, 82 classes) and test set (1,155 images, 11,550 captions, 20 classes).

Evaluation Metrics: Inception score (IS) [24] and Fréchet Inception distance (FID) [10] to evaluate our generative image autoencoder model quantitatively.

We have used four standard metrics for the performance evaluation of our LSTM-based text autoencoder. These metrics are BLEU [19], METEOR [1], ROUGEL [17], CIDEr [18] to quantify the similarity between the ground reference sentence and autoencoder generated sentence. The official MSCOCO caption evaluation scripts[2] are used for evaluation.

The following evaluation metrics are used to assess the performance of our mapping networks:

---

[2] https://github.com/tylin/coco-caption

1. **Human Score:** We conducted user studies for evaluation of our end-to-end text-to-image and image-to-text systems. We sampled 2000 samples for each of the text-to-image and image-to-text cases, assigned 5 different users (excluding any of the authors) to score the results. Generated samples are evaluated on a 4.0 scale, 1.0 being the lowest (worst case) and 4.0 being the highest (best case). Point 4.0 is awarded when the text description matches completely the image without any errors. Point 3.0 is awarded when the text description partially matches the image with minor errors. Point 2.0 is awarded when there is somewhat matching between text and image. Point 1.0 is awarded when the image and text are unrelated.

2. **Class Accuracy:** The CUB-200-2011 dataset has 150 train classes and 50 test classes and the Oxford-102 dataset has 82 and 20 train and test classes, respectively. Classes define different species for birds. Birds in the same class are from the same species, have identical features and descriptions. It is logical to expect that when one modality embedding maps into the other modality embedding, they are mapped to the same class. We first obtain the embeddings of one modality (say modality **A**) call it **true embeddings in modality A**. Then we obtain the embeddings of other modality (say modality **B**) and use them to feed into the **"Modality B-To-Modality A" mapping network** to obtain the **fake embeddings in modality A**. Then the cosine similarity between the true embeddings and fake embeddings are computed and we do an *argmax* to determine the class of fake embedding based on the highest cosine similarity scores. Then we calculate the class accuracy based on class labels of the true embeddings and predicted class labels of the fake embeddings.

5   Result and Analysis

The Inception score and FID of the StackGAN autoencoder model have been reported and compared with the baselines in Table 1. All the reported results are statistically significant. Results clearly illustrate the best performance by StackGAN autoencoder.

The scores for LSTM-based text autoencoder model are reported as: on CUB dataset (BLEU-1 :87.52, BLEU-4: 72.55, METEOR: 49.89, ROUGE-L : 89.04, CIDEr: 7.04) and on Oxford-102 (BLEU-1 :85.35, BLEU-4: 70.61, METEOR: 48.06, ROUGE-L : 86.69, CIDEr: 6.59)

The sample outputs of the End-To-End Image-To-Text synthesis are shown in Figure 4; Figure 3 depicts the sample outputs for Text-to-Image synthesis network. Evaluation metric scores are reported in Table 2 for GAN-based end-to-end networks and MMD-based networks.

The human score indicates that there exist some correlations between the two modalities of Image-to-Text and Text-to-Image synthesis systems. However, **Class Accuracy** score is indicative of a low semantic correlation between the two modalities for both GAN based and MMD based Image-to-Text and Text-to-Image systems. The scores at Table 2 imply similar performances by both the GAN-based system and the MMD-based system.

Dataset	Model	Inception Score ↑	FID ↓
CUB	**ResNet autoencoder with latent dimension** 1024	1.08 ± 0.03	275.71
CUB	**Com-Rec autoencoder with latent dimension** 1024	1.15 ± 0.02	248.95
CUB	**Compressive autoencoder with latent dimension** 1024	1.09 ± 0.08	283.64
CUB	**StackGAN autoencoder**	**3.66 ± 0.09**	**21.58**
Oxford-102	**StackGAN autoencoder**	**3.23 ± 0.14**	**57.96**

Table 1: Inception scores, Fréchet Inception distance (FID) (computed for 256×256 images)

Model	Metric	CUB-200-2011		Oxford-102	
		Image-To-Text	Text-To-Image	Image-To-Text	Text-To-Image
GAN-based	Human Score	**2.3**	**2.5**	**2.5**	**2.4**
	Class Accuracy*	**2.3**	**2.8**	**6.6**	**4.5**
MMD-based	Human Score	2.1	2.4	2.2	2.3
	Class Accuracy*	2.0	2.5	5.6	4.1

Table 2: Human Score, Class Accuracy of GAN-based and MMD-based Image-To-Text and Text-To-Image End-To-End systems on CUB and Oxford-102 datasets. (* means the values reported are in percentage).

this bird is smaller, short beaked, yellow superciliary has shades of light brown to dark brown on the wings

a tiny bird, with a tiny head and bill, brown coverts, white secondaries, a black back, grey outer retices, and a white breast

the build of the bird is chubby

the bird has a small black bill that is somewhat curved

Fig. 3: End-To-End Network Output for Birds: Text-To-Image Synthesis

the flowers has thin lavender petals

this small yellow flower has overlapping pointed around white pistils

the flower shown has smooth petals that are fused and separately arranged orangish in region

this flower has large purple petals petals, with pointed edges

Fig. 4: End-To-End Network Output for Flowers: Image-To-text Synthesis

6   Error Analysis

A thorough analysis revealed scenarios for possible reasons for errors. The strength of the end-to-end model is dependant on strength of the individual components, as models are trained separately; so error in each component has a cumulative impact on the end-to-end network. The autoencoders are trained separately without sharing weights or information, as a result, the latent spaces are disjoint, which further makes it harder for the mapping networks to align learned embedding spaces on the semantic level.

7   Conclusion

In this paper, we proposed a novel self-supervised approach that performs the cross-modal translation from image to text and text to image. We showed that our Text-to-Image and Image-to-Text synthesis networks learn to map the semantic space of one modality to the semantic space of the other modality in an unsupervised fashion. However, we figure out while learning the mapping, the semantic correlation across the modalities is low. Though the current end-to-end network depicts low cross-modal semantic alignments, as part of the future work, these learned network weights can be used as initialization for the synthesis networks and the entire network can be fine-tuned till the embeddings of the two modalities semantically align in the latent space.

Acknowledgments

References

1. Banerjee, S., & Lavie, A.: Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In: Proceedings of the acl workshop on intrinsic and extrinsic evaluation (2005).
2. Berlinet, A., & Thomas-Agnan, C.: Reproducing kernel Hilbert spaces in probability and statistics. In: Springer Science & Business Media (2011).
3. Chelba, C., Mikolov, T., Schuster, M., Ge, Q., Brants, T., Koehn, P., & Robinson, T.: One billion word benchmark for measuring progress in statistical language modeling. arXiv preprint arXiv:1312.3005 (2013).
4. Das, A., Kottur, S., Gupta, K., Singh, A., Yadav, D., Moura, J. M., Parikh, D., & Batra, D. (2017).: Visual dialog. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 326-335).
5. Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition (pp. 248-255).
6. Fan, L.: Semantic visual search engine. In: US Patent 7,865,492. Gehring, J., Auli, M., Grangier, D., Yarats, D., & Dauphin, Y. N. (2017). Convolutional sequence to sequence learning. arXiv preprint arXiv:1705.03122, (2011).
7. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014).: Generative adversarial nets. In: Advances in neural information processing systems (pp. 2672-2680).
8. Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B. & Smola, A. (2012a).: A kernel two-sample test. In: Journal of Machine Learning Research, 13, 723-773.
9. He, K., Zhang, X., Ren, S., & Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 770-778) (2016).
10. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., & Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: Advances in neural information processing systems (pp. 6626-6637) (2017).
11. Jiang, F., Tao, W., Liu, S., Ren, J., Guo, X., & Zhao, D.: An end-toend compression framework based on convolutional neural networks. In: IEEE Transactions on Circuits and Systems for Video Technology, 28, 3007-3018 (2017).
12. Kiros, R., Salakhutdinov, R., & Zemel, R.: Multimodal neural language models. In: International Conference on Machine Learning (pp. 595-603), (2014).
13. Lample, G., Conneau, A., Denoyer, L., & Ranzato, M.: Unsupervised machine translation using monolingual corpora only. arXiv preprint arXiv:1711.00043, (2017).
14. Lample, G., Ott, M., Conneau, A., Denoyer, L., & Ranzato, M.: Phrase-based & neural unsupervised machine translation. arXiv preprint arXiv:1804.07755, (2018).
15. Lao, Q., Havaei, M., Pesaranghader, A., Dutil, F., Jorio, L. D., & Fevens, T.: Dual adversarial inference for text-to-image synthesis. In: Proceedings of the IEEE International Conference on Computer Vision (pp. 7567-7576), (2019).
16. Li, C.-L., Chang, W.-C., Cheng, Y., Yang, Y., & Póczos, B.: Mmd gan: Towards deeper understanding of moment matching network. In: Advances in Neural Information Processing Systems (pp. 2203-2213), (2017).
17. Lin, C.-Y. (2004). ROUGE: A package for automatic evaluation of summaries. In Text Summarization Branches Out (pp. 74-81). In: Barcelona, Spain: Association for Computational Linguistics. https://www.aclweb.org/anthology/W04-1013.
18. Vedantam R.,Lawrence Zitnick C. & Parikh D.: Cider: Consensusbased image description evaluation. In: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 4566– 4575.

19. Papineni, K., Roukos, S., Ward, T., & Zhu, W.-J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th annual meeting on association for computational linguistics (pp. 311-318). Association for Computational Linguistics, (2002).
20. Pedersoli, M., Lucas, T., Schmid, C., & Verbeek, J.: Areas of attention for image captioning. In: Proceedings of the IEEE International Conference on Computer Vision (pp. 1242-1250), (2017).
21. Pennington, J., Socher, R., & Manning, C. D.: Glove: Global vectors for word representation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP) (pp. 1532-1543) (2014).
22. Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., & Lee, H.: Generative adversarial text to image synthesis. arXiv preprint arXiv:1605.05396, (2016a).
23. Reed, S. E., Akata, Z., Mohan, S., Tenka, S., Schiele, B., & Lee, H.: Learning what and where to draw. In: Advances in neural information processing systems (pp. 217-225), (2016b).
24. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., & Chen, X.: Improved techniques for training gans. In: Advances in neural information processing systems (pp. 2234-2242), (2016).
25. Theis, L., Shi, W., Cunningham, A., & Huszár, F.: Lossy image compression with compressive autoencoders. arXiv preprint arXiv:1703.00395, (2017).
26. Venugopalan, S., Anne Hendricks, L., Rohrbach, M., Mooney, R., Darrell, T., & Saenko, K.: Captioning images with diverse objects. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 5753-5761), (2017).
27. Vinyals, O., Toshev, A., Bengio, S., & Erhan, D.: Show and tell: A neural image caption generator. In: Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 3156-3164), (2015).
28. Wah, C., Branson, S., Welinder, P., Perona, P., & Belongie, S.: The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001 California Institute of Technology, (2011).
29. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., & Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: International conference on machine learning (pp. 2048-2057), (2015).
30. Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., & Metaxas, D. N. :Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision (pp. 5907-5915), (2017).
31. Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., & Metaxas, D. N. : Stackgan++: Realistic image synthesis with stacked generative adversarial networks.In: IEEE transactions on pattern analysis and machine intelligence, 41, 1947-1962, (2018).
32. Tolstikhin, I. O., Sriperumbudur, B. K., & Schölkopf, B.: Minimax estimation of maximum mean discrepancy with radial kernels. In: Advances in Neural Information Processing Systems (pp. 1930-1938), (2016).