

This ICCV 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.

CIGLI: Conditional Image Generation from Language & Image

Xiaopeng Lu*

Lynnette Ng*

Jared Fernandez* Hao Zhu*

Carnegie Mellon University

xiaopen2@cs.cmu.edu, {lyn

{lynnetteng, jaredfern, zhuhao}@cmu.edu

Abstract

Multi-modal generation has been widely explored in recent years. Current research directions involve generating text based on an image or vice versa. In this paper, we propose a new task called CIGLI: Conditional Image Generation from Language and Image. Instead of generating an image based on text as in text-image generation, this task requires the generation of an image from a textual description and an image prompt. We designed a new dataset to ensure that the text description describes information from both images, and that solely analyzing the description is insufficient to generate an image. We then propose a novel language-image fusion model which improves the performance over two established baseline methods, as evaluated by quantitative (automatic) and qualitative (human) evaluations. The code and dataset is available at https://github.com/vincentlux/CIGLI.

1. Introduction

Multimodal reasoning over visual and language is a long-standing research problem with the ultimate goal of building a system to connect and understand information across two vastly different inputs. Current research in this area primarily focuses on classification tasks, where the model jointly understands the text and image information to produce a classification label. State-of-the-art multimodal models separately model the caption and the images before jointly processing the combined representation to perform a *TRUE/FALSE* evaluation [16, 13, 2].

One such discriminative multimodal reasoning task is Natural Language Visual Reasoning (NLVR) [14] which defines a binary classification task to test model's reasoning capabilities across modalities. Given a sentence description and three synthetic images, the task is to predict whether the sentence correctly described the images or not. In NLVR2 [15], realistic natural images are added though the task remains the same.

While advances in modeling have led to improved results on these tasks, current models are largely based around BERT-style autoencoding architectures which are not wellsuited for generation. As opposed to classification tasks which only require models to output a label, image generation from natural language captions presents further challenges in the multimodal regime as it requires a much richer output than a single discriminative label [1, 4, 10]. In this work, we propose a novel generative task: given a description describing two images, and one of the images, can a model learn to generate the second image that is semanti*cally* correct? For example, given a text description "There are two dogs in total", and a first image with one dog in it, the model should be able to generate the second image with exactly one dog. While recent work has explored conditional generation of images from natural language [9], to our knowledge there has been no prior work constraining generation based on both language and visual inputs. Our task provides the model with the caption and the first image in order to generate the second image, requiring the examination of both the language and visual prompts and understand the semantic relationships between the input text and image. This task requires the model to both understand the natural language caption and identify the missing visual components from the first image to generate the second image in the pair. Thus, it cannot be performed by conditioning solely on text as the contents of the second image are directly dependent on the contents of the first.

We construct a new dataset based off the NLVR2 task [15] and perform image generation on two GAN-based architectures. We further evaluate their performance using both quantitative (automatic) and qualitative (human) evaluations. While our results are preliminary, we hope to inspire further directions in language-visual reasoning, especially in the combined understanding of both mediums.

2. Related Work

Many neural approaches to text-to-image generation rely on architectures built around Generative Adversarial Network (GANs). Several GAN-based models composed im-

^{*} Equal Contribution

ages by iteratively retrieving word vectors for different portions of the image and up-scaling the image resolution [18, 19]. Alternate GAN-based approaches attempted different combinations of inputs to the generators and discriminators, such as: text features inputs into both generator and discriminator [11], text and image features input both the generator and discriminator [20], or text features into the generator and image into the discriminator [17]. Recent text-to-image models have also been developed around the transformer architectures such as X-LXMERT [3] and DALL-E [9].

3. Dataset Construction

We construct a new dataset by combining the train and dev splits of the original NLVR2 dataset and filtering for image pairs and captions which require visual reasoning. Each data point in the source NVLR2 dataset contains two images and a natural language sentence. Next, to ensure that the generation task requires reasoning about both the language and visual prompt, we removed all *FALSE* examples, as negative captions can trivially be satisfied by outputting noise for the generated image.

Additionally, we manually examined 600 randomly sampled images to determine a set of heuristic patterns found in the text captions that require reasoning about both images. These patterns consist of captions where the model cannot perform image generation based on the caption alone; it must also analyze the first image to determine which part of the caption it satisfies. Figure 1 shows examples of filtered data. The final dataset consists of 24,178 valid data points that were found to satisfy our heuristic-based filtering ¹.

4. Methodology

In this section, we outline our methodology for image generation baselines task with our constructed dataset. We first perform an evaluation using language-only baseline models from available pre-trained AttnGAN and DF-GAN models, then select the better of the two for further improvements. We improve the model through fine-tuning the model and fusion of image features. The two model architectures are reflected in Figure 2.

4.1. Baseline Method

We use a language-only model for our baseline image generation. The model is fed an initial image, named the "First Image", and the original caption that describes both images. From the first image, we generate a caption for it through an image to caption algorithm. We then concatenate the original caption with the generated caption, and



Figure 1: Cases of qualified and unqualified data. (a) is qualified because counting about the number of objects across both images. (b) is qualified because it requires distinguishing which textual condition is satisfied by the first image before generating the second image. (c) is unqualified because the second image can be generated without looking at the first. (d) is unqualified because the model can ignore the first image and generate a single dog.



Figure 2: Baseline model (upper image) vs. Our model (lower image)

pass it through a text-to-image-generation algorithm to produce the final generated image.

In the first step, image captioning [7] translates images into text in two steps: (1) image encoding; (2) language generation. We use the following image captioning model: LSTM [6] with features from pretrained ResNet50 [5], which achieves 35.9 BLEU@4 on MSCOCO [8]. We use beam search (beamSize=10) to get the best captions for each image. We generate 10 captions for each image through this algorithm.

We then concatenate all ten of the generated captions together with the original NLVR caption in order to get as much information for the text-to-image generation step. This becomes the input for the image generation model. We run image generation on two pre-trained models: At-tnGAN [18] and DF-GAN [17]. Both models were worked on based on a pre-trained MSCOCO model. Although these

lfiltering code is available at https://github.com/ vincentlux/CIGLI/blob/main/filter.py

models have better output with other datasets or subsets of MSCOCO, we chose to use the entire dataset as it has more diverse objects, in line with our natural image generation task. AttnGAN was run with a text embedding vector of size 256, and a training batch size of 100, and produces an image of a 256 * 256 resolution.

4.2. Proposed Method

We perform automatic evaluations on baseline generated images and improved on DF-GAN, as it is the better of the two models. We use a language-and-vision model for generation (Figure 2, lower image). We represent the First Image as a vector from ResNet50 features pretrained on MSCOCO. We represent the text caption as an LSTM vector. We then concatenate both image and language vectors, and input the combined representation into an image generation model. We use DF-GAN as our image generation model to provide direct comparison with our baseline model. We employ two fusion techniques: (1) concatenation of the image and text feature vectors and (2) sum of the image and text feature vectors.

We train our proposed network end-to-end without freezing the text or image encoder sub-networks, with 120 epochs and a batch size of 24.

4.3. Evaluations

We employ automatic evaluation using the UNITER metric and the Inception Score metric, and human evaluation. We used automatic evaluation to evaluate our baseline models of AttnGAN and DF-GAN, before picking the better model to improve on; then used both automatic and human evaluation to evaluate our proposed model architecture.

4.3.1 Automatic Evaluation

UNITER metric To provide an automated evaluation of the generated image quality, we sought to determine if the generated images were sufficient for a fine-tuned UNITER model [2] to recognize that the pair of images satisfied the reasoning statement in the caption.

We performed visual feature extraction using a Faster RCNN network to extract 768-dimensional visual features from both the generated images and the source images. Prior to use as a metric, the UNITER model was finetuned on the NLVR2 task for 8000 steps. The UNITER model was then provided 1805 true examples in which the caption correctly describes the image pairs. To compare against existing generative baselines, we swapped out the second image in the pair with the generation in the data point.

Inception Score To understand the quantitative quality of the image, we use the Inception Score [12]. Although the quality of image speaking of naturalness and diversity

is not our main concern, inception score serves to evaluate the general quality of the generated image.

4.3.2 Human Evaluation

Quantitative intrinsic metrics fail to capture the semantic correctness of the generated image. Therefore, human evaluation plays a significant role in our experiment. Based on previous research [9], we define three dimensions for human evaluation:

Naturalness: the evaluators were asked to use their intuition on how natural the generated image looks

Relevance whether the images are relevant to the caption **Correctness**: whether the generated image is semantically correct (eg. number of objects, object type etc)

We constructed a web application for human evaluation. For each data point, the evaluators are presented with the caption and the First Image. They are asked if the image satisfies the caption. If they select "NO", they are presented with the generated images to evaluate further. Evaluators are presented with the generated images one by one and asked to provide a YES/NO binary rating to each of the metric. This is a blind evaluation; the evaluators do not know which model generated which image. We did not show the model name to the evaluators to alleviate the bias related to knowing how the image is generated. Figure in the appendix shows an example of a data point on the application presented to the evaluators for each type of rating scheme.

5. Results

The generated images are presented in Table 4. The automated evaluations using UNITER and Inception Score results are in Table 1.

For human evaluation, we used four evaluators. The results are presented in Figure 3. The evaluators identified that in 34% of the data points presented, the first image had already satisfied the caption. We then evaluated the remaining 66% of images for their naturalness, relevance and accuracy. We evaluated the percentage of positive annotations for each metric; i.e. the percentage where evaluators annotate YES to the metric. As expected, the NLVR2 image ranks at 100% for all three categories. The concatenated fusion model (DF-GAN concatenated) produces higher human ratings as compared to the summation fusion model (DF-GAN sum). These results are consistent with the automated evaluation, where DF-GAN Concatenated image reports the best result quality.

6. Discussion

Previous work in image generation typically generate an output image from a single text caption. In this study, our image generation model is conditioned on both an input image to and a text caption. Our proposed DF-GAN

	Inception Score↑	UNITER-Accuracy [↑]	UNITER-Consistency↑
Original Image	-	76.07	64.96
Baseline Models			
AttnGAN	5.50 ± 0.43	41.55	25.13
DF-GAN	9.59 ± 0.59	42.94	26.25
Proposed Models			
DF-GAN Concatenated	11.30 ± 0.83	46.54	28.50
DF-GAN Sum	$9.67 {\pm} 0.53$	44.97	28.80

Table 1: Automatic Evaluation Score Comparison



Figure 3: Human evaluation results

concatenated fusion model outperforms the sum-fusion and the baseline models, suggesting that the images generated from this architecture are more similar to the original image than other architectures. Still, the image generation output has a long way to go if we compare the score with original NLVR2 original images, showing that CIGLI task is far from been solved, and more future work can be done on this task to improve the performance.

Automatic evaluation score analysis The accuracy of the UNITER model on the ground truth original image serves as an upper bound baseline for the performance of generated images. The DF-GAN network trained with our proposed technique outperforms both language-only AttnGAN and DF-GAN baselines on both accuracy and consistency. This suggests that the images generated with our proposed metric are better able to reflect the necessary visual features needed for the UNITER model to recognize that the image pair is a valid reasoning statement.

Inception score reflects the naturalness and diversity of the generated images. Although this metric cannot measure the semantic correctness of the image, it can still be used as a way to help understand how good is the generation quality in general. From our results, our proposed concatenated model improves the Inception score by 1.71 compared with the DF-GAN baseline. We posit this is due to the extra NLVR2 training data we used for conditional image generation.

Human evaluation score analysis The human evaluation results shows promise as our proposed methods outperforms the baseline method. Despite the lack of naturalness and correctness of our proposed models as compared to the original NLVR2 images, a manual inspection observes that our method performs pretty well on common objects such as dogs, and not so well on uncommon objects like alpacas. Captions that tend to go along with good output images are quantitative captions (eg "there are two dogs") describing a single object, rather than a comparison of images (eg "the right image contains... while the left image contains...") or captions describing multiple objects in an image (eg "there is a girl using a smartphone").

7. Conclusions

In this paper, we introduce CIGLI, a new task for conditional image generation from both text and images. This task poses new challenges to vision-language research, as it requires model to perform semantic reasoning across both modalities before image generation. We also propose a new image-text fusion model based on DF-GAN, which improves the performance compared with two baseline models. We hope that this new dataset and the corresponding model can be the start point of future research in conditional image generation from language and image.

References

- Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *International Conference on Machine Learning*, pages 1691–1703. PMLR, 2020.
- [2] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European Conference on Computer Vision*, pages 104–120. Springer, 2020.
- [3] Jaemin Cho, Jiasen Lu, Dustin Schwenk, Hannaneh Hajishirzi, and Aniruddha Kembhavi. X-lxmert: Paint, caption and answer questions with multi-modal transformers, 2020.



Figure 4: Qualitative results

- [4] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Rezende, and Daan Wierstra. Draw: A recurrent neural network for image generation. In Francis Bach and David Blei, editors, Proceedings of the 32nd International Conference on Machine Learning, volume 37 of Proceedings of Machine Learning Research, pages 1462–1471, Lille, France, 07–09 Jul 2015. PMLR.
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [6] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [7] Ryan Kiros, Ruslan Salakhutdinov, and Rich Zemel. Multimodal neural language models. In *International conference* on machine learning, pages 595–603. PMLR, 2014.
- [8] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [9] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. arXiv preprint arXiv:2102.12092, 2021.
- [10] Scott Reed, Zeynep Akata, Santosh Mohan, Samuel Tenka, Bernt Schiele, and Honglak Lee. Learning what and where to draw. In Advances in Neural Information Processing Systems, 2016.
- [11] Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1060– 1069, New York, New York, USA, 20–22 Jun 2016. PMLR.

- [12] Tim Salimans, Ian J Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NIPS*, 2016.
- [13] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vl-bert: Pre-training of generic visuallinguistic representations. In *International Conference on Learning Representations*, 2019.
- [14] Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. A corpus of natural language for visual reasoning. In *Proceedings* of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pages 217–223, 2017.
- [15] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. arXiv preprint arXiv:1811.00491, 2018.
- [16] Hao Tan and Mohit Bansal. Lxmert: Learning crossmodality encoder representations from transformers. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 5103–5114, 2019.
- [17] Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Xiao-Yuan Jing, Fei Wu, and Bingkun Bao. Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis, 2021.
- [18] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Finegrained text to image generation with attentional generative adversarial networks. In CVPR, 2018.
- [19] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*, 2017.
- [20] Yizhe Zhu, Mohamed Elhoseiny, Bingchen Liu, Xi Peng, and Ahmed Elgammal. A generative adversarial approach for zero-shot learning from noisy texts. In *CVPR*, 2018.