

Leveraging Visual Question Answering to Improve Text-to-Image Synthesis

Stanislav Frolov^{1,2}, Shailza Jolly^{1,2}, Jörn Hees², Andreas Dengel^{1,2}

¹ Technical University of Kaiserslautern, Germany

² German Research Center for Artificial Intelligence (DFKI), Germany

firstname.lastname@dfki.de

Abstract

Generating images from textual descriptions has recently attracted a lot of interest. While current models can generate photo-realistic images of individual objects such as birds and human faces, synthesising images with multiple objects is still very difficult. In this paper, we propose an effective way to combine Text-to-Image (T2I) synthesis with Visual Question Answering (VQA) to improve the image quality and image-text alignment of generated images by leveraging the VQA 2.0 dataset. We create additional training samples by concatenating question and answer (QA) pairs and employ a standard VQA model to provide the T2I model with an auxiliary learning signal. We encourage images generated from QA pairs to look realistic and additionally minimize an external VQA loss. Our method lowers the FID from 27.84 to 25.38 and increases the R-prec. from 83.82% to 84.79% when compared to the baseline, which indicates that T2I synthesis can successfully be improved using a standard VQA model.




Figure 1: We use concatenated question answer (QA) pairs from VQA data as additional input samples for T2I training. A standard VQA model can then be used to provide an additional loss to the T2I model.

1 Introduction

Text-to-image synthesis (T2I), the task to generate realistic images given textual descriptions, received a lot of attention in recent years (Reed et al., 2016; Zhang et al., 2016; Zhang et al., 2017; Xu et al., 2017; Zhu et al., 2019; Li et al., 2019). T2I synthesis can be seen as the inverse of image captioning. Given a caption, the model is trained to produce realistic images that correctly reflect the meaning of the input captions. Many existing T2I methods use Generative Adversarial Networks (GANs) (Goodfellow et al., 2014). GANs consist of two artificial neural networks that play a game in which a discriminator is trained to distinguish between real and generated images, while a generator is trained to produce images to fool the discriminator. They have successfully been applied to many image synthesis applications such as image-to-image translation (Isola et al., 2016; Zhu et al., 2017), image super-resolution (Ledig et al., 2016), and image in-painting (Yeh et al., 2016).

Similarly, Visual Question Answering (VQA) (Antol et al., 2015) emerged as an important task to build systems that better understand the relationship between vision and language by learning to answer

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.




Figure 2: An overview of our architecture with our VQA extension highlighted in red. Given text inputs t (captions) and t_{QA} (concatenated QA pairs) we generate images \hat{x}_g and \tilde{x}_g . As in AttnGAN, we use the discriminator and DAMSM model to differentiate between real (omitted in the figure for brevity) and generated images as well as image-text pairs, leading to loss \mathcal{L}_{ADV} , and DAMSM loss \mathcal{L}_{DAMSM} for improved image-text alignment. Images generated from QA pairs are passed through a VQA model resulting in the additional loss \mathcal{L}_{VQA} .

questions about an image. It can be seen as image conditioned question answering, which encourages the model to both look at the image and understand the question to predict the correct answer.

A good T2I model should produce images that look realistic and correctly reflect the semantic meaning of the input description. Considering the complexity of natural language descriptions and difficulty to produce photographic images, current models struggle to achieve these goals. In this paper we propose a simple, yet effective way to combine T2I with VQA to improve both the image quality and image-text alignment of generated images of a T2I model by leveraging the questions and answers (QA) provided in the VQA 2.0 dataset (Antol et al., 2015). Both captions and QA pairs can describe the overall image or very specific details. In fact, many QA pairs can be rephrased as captions and vice versa which further motivates to leverage VQA for T2I. Additionally, the VQA 2.0 dataset contains complementary images with different answers for the same question which requires the T2I model to learn to pay close attention to the input in order to generate an image that correctly reflects the meaning of the input text. To leverage the VQA 2.0 dataset for T2I, we concatenate QA pairs and use them as additional training samples in our T2I pipeline. Images generated from QA pairs can subsequently be used as inputs to a VQA model which can provide an additional learning signal to the T2I generator. See Figure 1 for an overview of our approach.

2 Related Work

Initial T2I approaches (Reed et al., 2016; Dash et al., 2017) adopted the conditional GAN (cGAN) (Mirza and Osindero, 2014) and AC-GAN (Odena et al., 2016) ideas to replace the conditioning variable by a text embedding which allows to condition the generator on a textual description. Analog to many current approaches, our approach is also based on AttnGAN (Xu et al., 2017), which incorporates an attention mechanism on word features to allow the network to synthesize fine-grained details.

In terms of using VQA for T2I, to our knowledge the only other approach is VQA-GAN (Niu et al., 2020). However, in contrast to their approach, our architecture is simpler, we do not use layout information, work on individual QA pairs, and produce higher resolution images (256x256 vs. 128x128).

3 Approach

We extend the AttnGAN (Xu et al., 2017) architecture to leverage question and answer (QA) pairs from the VQA 2.0 (Goyal et al., 2017) dataset by appending a VQA model (Kazemi and Elqursh, 2017). In addition to generating images from image descriptions, our model is also trained to produce images given a QA pair and minimize an external VQA loss. See Figure 2 for an overview. After revisiting the individual components, we explain our extension in more detail.

3.1 T2I: AttnGAN

AttnGAN (Xu et al., 2017) consists of a multi-stage refinement pipeline that employs attention-driven generators for fine-grained T2I synthesis. A pre-trained bidirectional LSTM (BiLSTM) (Schuster and

Paliwal, 1997) is used to extract global sentence as well as individual word features. Discriminators jointly approximate the conditional and unconditional distributions simultaneously. Additionally, they use an image-text matching loss at the word level called Deep Attentional Multimodal Similarity Model (DAMSM) is employed to guide the image generation process. The attention-driven generator together with the DAMSM loss help the generator to be able to focus on individual words to synthesize fine-grained details and improve the semantic alignment between input description and final image.

3.2 VQA Model

We use a very basic and widely used VQA model, proposed in (Kazemi and Elqursh, 2017)¹. Given an image I and question q , the model uses image features extracted from a pre-trained ResNet (He et al., 2016), an LSTM (Hochreiter and Schmidhuber, 1997) based question embedding, and employs stacked attention (Yang et al., 2016) to produce probabilities over a fixed set of answers. The VQA loss, given in Equation 1, is simply an average over the negative log-likelihoods over all the correct answers a_1, a_2, \dots, a_K .

$$\mathcal{L}_{\text{VQA}} = \frac{1}{K} \sum_{k=1}^K -\log P(a_k | I, q) \quad (1)$$

3.3 T2I + VQA

We extend AttnGAN (Xu et al., 2017) by appending the VQA model (Kazemi and Elqursh, 2017), and create additional training samples by concatenating question and answer (QA) pairs. Given captions t and QA pairs t_{QA} , we generate fake images \hat{x}_g , and \tilde{x}_g . Real images are denoted as x_r . Next, the VQA model takes the image generated from the QA pair and corresponding question to produce an answer. Similar to (Niu et al., 2020), we use the VQA loss to predict the correct answer to guide the image generator. At the same time, the discriminators encourage the generated images to be realistic and matching to their corresponding text input. The DAMSM loss (Xu et al., 2017) further helps to focus on the individual words to generate fine-grained details. The final objective of our generator is defined in Equation 2, where \mathcal{L}_{ADV} is the adversarial (conditional and unconditional) loss, and $\mathcal{L}_{\text{DAMSM}}$ is the DAMSM loss, both as described in (Xu et al., 2017). The adversarial loss and DAMSM loss are applied to both captions and QA pairs and their correspondingly generated images. The VQA loss \mathcal{L}_{VQA} is applied only to images generated from QA pairs.

$$\begin{aligned} \mathcal{L}_G &= \mathcal{L}_{\text{ADV}} + \mathcal{L}_{\text{DAMSM}}(\hat{x}_g, t) + \mathcal{L}_{\text{DAMSM}}(\tilde{x}_g, t_{\text{QA}}) + \mathcal{L}_{\text{VQA}}(\tilde{x}_g, q) \\ \mathcal{L}_{\text{ADV}} &= \underbrace{-\mathbb{E}[\log D(\hat{x}_g)] - \mathbb{E}[\log D(\tilde{x}_g)]}_{\text{unconditional loss}} - \underbrace{\mathbb{E}[\log D(\hat{x}_g, t)] - \mathbb{E}[\log D(\tilde{x}_g, t_{\text{QA}})]}_{\text{conditional loss}} \end{aligned} \quad (2)$$

The discriminators are trained to classify between real and fake images, as well as image-caption and image-QA pairs to simultaneously approximate the conditional and unconditional distributions by minimizing the modified loss defined in Equation 3.

$$\begin{aligned} \mathcal{L}_D &= \underbrace{-\mathbb{E}[\log D(x_r)] - \mathbb{E}[\log(1 - D(\hat{x}_g))] - \mathbb{E}[\log(1 - D(\tilde{x}_g))]}_{\text{unconditional loss}} \\ &\quad - \underbrace{\mathbb{E}[\log D(x_r, t)] - \mathbb{E}[\log(1 - D(\hat{x}_g, t))] - \mathbb{E}[\log(1 - D(\tilde{x}_g, t_{\text{QA}}))]}_{\text{conditional loss}} \end{aligned} \quad (3)$$

4 Experiments

We train three different variants of our extension. The first two are naive extensions in which we do not change the discriminator loss of AttnGAN. Instead, we simply append the VQA model, sample a QA

¹<https://github.com/Cyanogenoid/pytorch-vqa>

pair during training and add the external VQA loss to the overall generator loss function. We experiment with an end-to-end training approach where we start with a randomly initialized VQA model, and a pre-trained VQA model. Next, we train a model in which we change the discriminator and generator loss functions. In other words, given a QA pair, the model is not only trained to minimize the VQA loss for that particular QA pair, but also to produce realistic and matched images as judged by the discriminator and DAMSM loss. For fair comparison we re-train and re-evaluate AttnGAN using the codebase provided by the authors².

4.1 Datasets and Evaluation

We use the commonly used COCO (Lin et al., 2014) and VQA 2.0 (Goyal et al., 2017) datasets to train our model. COCO depicts complex scenes and multiple interacting objects and contains around 80k images for training and 40k images for testing. Each image has five captions. VQA 2.0 is a large dataset of question answer pairs based on the COCO images with roughly 400k QAs for training and 200k for testing, hence extensively used by VQA researchers (Tan and Bansal, 2019; Kim et al., 2018; Lu et al., 2019). It contains complementary images for the same question, such that the model learns to look closely at the image before answering instead of deploying language biases (Antol et al., 2015).

We evaluate the quality and diversity of generated images using the Inception Score (IS) (Salimans et al., 2016) and Fréchet Inception Distance (FID) (Heusel et al., 2017). To compute the IS³ and FID⁴ we generate 30k images from 30k randomly sampled test captions. R-prec. (Xu et al., 2017) is used to evaluate the semantic alignment between generated images and input captions. Although R-prec. might be unreliable, as current models seem to achieve higher scores than real images (Hinz et al., 2019), we include it for reference since it is still commonly used. Similar to (Niu et al., 2020), we evaluate the VQA accuracy of our models by generating images from test QA pairs and passing them to the pre-trained VQA model with the corresponding question and answer.

4.2 Results

Method	IS \uparrow	FID \downarrow	R-prec. \uparrow	VQA Acc. \uparrow
Real Images	34.88	6.09	68.58	60.00
AttnGAN (Xu et al., 2017)	26.66	27.84	83.82	43.00
AttnGAN + end-to-end VQA	25.22	30.68	82.68	42.85
AttnGAN + pre-trained VQA	26.02	28.72	84.25	42.83
AttnGAN + pre-trained VQA + adapted loss	<u>26.64</u>	25.38	84.79	43.75

Table 1: Results on the COCO test dataset. First row contains scores for real images (excluding VQA Acc.) as reported in (Hinz et al., 2019). We re-train and re-evaluate the baseline AttnGAN (second row). Third and fourth row are our naive extensions in which we simply append a VQA model for an external VQA loss for images generated from QA pairs. In the last row, we change the discriminator and generator losses to also encourage images generated from QA pairs to look realistic and match the input (similar to standard AttnGAN losses for images generated from captions). We train each model for 120 epochs, select the checkpoint with the best IS and report corresponding FID, R-prec., and VQA accuracy.

As can be seen in Table 1, merely adding the external VQA loss to the generator impairs the performance, regardless of using a pre-trained VQA model or training it as part of the pipeline in an end-to-end way. We hypothesize this is due to the images produced from QA pairs not being encouraged to look realistic during training and the generator struggling to minimize the external VQA loss for images generated from QA pairs, on the one hand, and standard AttnGAN losses for images generated from captions, on the other hand.

²<https://github.com/taoxugit/AttnGAN>

³<https://github.com/sbarratt/inception-score-pytorch>

⁴<https://github.com/mseitzer/pytorch-fid>

Therefore, in our third experiment (last row) we change the overall objective and encourage all generated images to be realistic and matched to the input description using the conditional and unconditional losses from the discriminator and DAMSM loss (as in standard AttnGAN). While achieving almost identical IS, our model greatly improves the FID from 27.84 to 25.38, which indicates better image quality and diversity. Additionally, the images produced by our extension are better aligned to the input descriptions as indicated by the improvement of R-prec. from 83.82 to 84.79, and VQA accuracy from 43.00 to 43.75. Since the VQA 2.0 dataset contains complementary image and QA pairs, the slight variation in linguistic inputs might hence help the model to generate better images. Our results show that it is possible to improve T2I synthesis by simply appending a standard pre-trained VQA model, leveraging the VQA 2.0 dataset as additional supervision and encouraging the model to also produce realistic images from QA pairs. Although we show that already a simple VQA model can help to improve the final T2I performance, we hypothesize that using a state-of-the-art VQA model might further improve the results.

5 Conclusion

In this paper we proposed a simple method to leverage VQA data for T2I via a combination of AttnGAN, a well-known T2I model, and a standard VQA model. By concatenating question answer (QA) pairs from the VQA 2.0 dataset we created additional training samples. Images generated from the QA pairs are passed to the VQA model which provides an additional learning signal to the generator. Our results show a substantial improvement over the baseline in terms of FID, but requires additional supervision. Possible future research directions could be to investigate whether our extension also boosts other T2I models, and the influence of more training data and the external VQA loss separately.

Acknowledgements

This work was supported by the BMBF project DeFuseNN (Grant 01IW17002) and the TU Kaiserslautern PhD program.

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *ICCV*.
- Ayushman Dash, John Cristian Borges Gamboa, Sheraz Ahmed, Marcus Liwicki, and Muhammad Zeshan Afzal. 2017. Tac-gan - text conditioned auxiliary classifier generative adversarial network. *ArXiv*, abs/1703.06412.
- Ian J. Goodfellow, Jean Pouget-Abadie, M. Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *NIPS*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *CVPR*.
- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*.
- Tobias Hinz, Stefan Heinrich, and Stefan Wermter. 2019. Semantic object accuracy for generative text-to-image synthesis. *ArXiv*, abs/1910.13321.
- S. Hochreiter and J. Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9:1735–1780.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 2016. Image-to-image translation with conditional adversarial networks. In *CVPR*.
- Vahid Kazemi and Ali Elqursh. 2017. Show, ask, attend, and answer: A strong baseline for visual question answering. *ArXiv*, abs/1704.03162.
- Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. 2018. Bilinear attention networks. In *NeurIPS*.

- Christian Ledig, Lucas Theis, Ferenc Huszár, José Antonio Caballero, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. 2016. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*.
- Wenbo Li, Pengchuan Zhang, Lei Zhang, Qiuyuan Huang, Xiaodong He, Siwei Lyu, and Jianfeng Gao. 2019. Object-driven text-to-image synthesis via adversarial training. In *CVPR*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *ECCV*.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*.
- Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets.
- Tianrui Niu, Fangxiang Feng, Lingxuan Li, and Xiaojie Wang. 2020. Image synthesis from locally related texts. In *ICMR*.
- Augustus Odena, Christopher Olah, and Jonathon Shlens. 2016. Conditional image synthesis with auxiliary classifier gans. In *ICML*.
- Scott E. Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. 2016. Generative adversarial text to image synthesis. In *ICML*.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved techniques for training gans. In *NIPS*.
- Mike Schuster and K. Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.*, 45:2673–2681.
- Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*.
- Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. 2017. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *CVPR*.
- Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. 2016. Stacked attention networks for image question answering. In *CVPR*.
- Raymond A. Yeh, Chen Chen, Teck-Yian Lim, Alexander G. Schwing, Mark Hasegawa-Johnson, and Minh N. Do. 2016. Semantic image inpainting with deep generative models. In *CVPR*.
- Han Zhang, Tao Xu, and Hongsheng Li. 2016. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*.
- Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N. Metaxas. 2017. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41:1947–1962.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*.
- Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. 2019. Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis. In *CVPR*.