

# Text as Neural Operator: Image Manipulation by Text Instruction

Tianhao Zhang\*  
Google Research  
bryanzhang@google.com

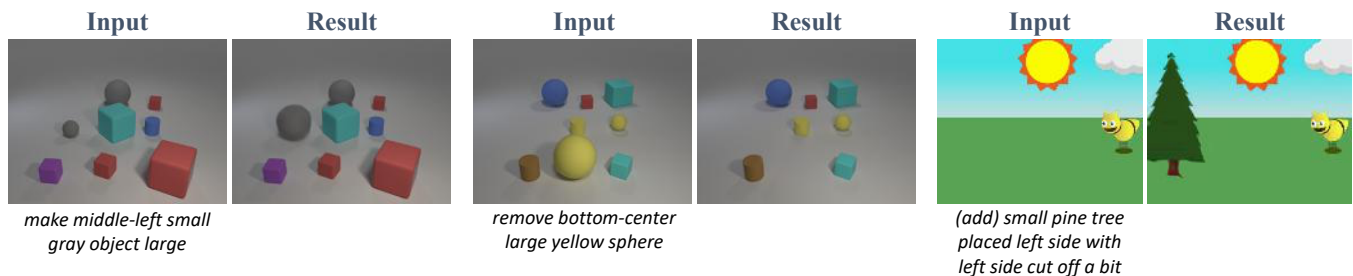
Hung-Yu Tseng†  
University of California, Merced  
htseng6@ucmerced.edu

Lu Jiang  
Google Research  
Carnegie Mellon University  
lujiang@google.com

Weilong Yang  
Waymo  
weilongyang@google.com

Honglak Lee  
University of Michigan  
honglak@eecs.umich.edu

Irfan Essa  
Google Research  
Georgia Institute of Technology  
irfanessa@google.com



**Figure 1: Image manipulation by text instruction. The input is multimodal consisting of a reference image and a text instruction. The results are synthesized images by our model.**

## ABSTRACT

In recent years, text-guided image manipulation has gained increasing attention in the multimedia and computer vision community. The input to conditional image generation has evolved from image-only to multimodality. In this paper, we study a setting that allows users to edit an image with multiple objects using complex text instructions to add, remove, or change the objects. The inputs of the task are multimodal including (1) a reference image and (2) an instruction in natural language that describes desired modifications to the image. We propose a GAN-based method to tackle this problem. The key idea is to treat text as neural operators to locally modify the image feature. We show that the proposed model performs favorably against recent strong baselines on three public datasets. Specifically, it generates images of greater fidelity and semantic relevance, and when used as a image query, leads to better retrieval performance.

## CCS CONCEPTS

• **Information systems** → **Multimedia content creation; Multimedia and multimodal retrieval**; • **Computing methodologies** → **Computer vision**.

\*Work done as a Google AI Resident.

†Work done during HY's internship at Google Research.



This work is licensed under a Creative Commons Attribution International 4.0 License.

MM '21, October 20–24, 2021, Virtual Event, China.  
© 2021 Copyright held by the owner/author(s).  
ACM ISBN 978-1-4503-8651-7/21/10.  
<https://doi.org/10.1145/3474085.3475343>

## KEYWORDS

Deep Neural Networks, Generative Adversarial Network, Multimodal Content Creation, Natural Language, Multimedia Retrieval

### ACM Reference Format:

Tianhao Zhang, Hung-Yu Tseng, Lu Jiang, Weilong Yang, Honglak Lee, and Irfan Essa. 2021. Text as Neural Operator: Image Manipulation by Text Instruction. In *Proceedings of the 29th ACM International Conference on Multimedia (MM '21)*, October 20–24, 2021, Virtual Event, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3474085.3475343>

## 1 INTRODUCTION

Image synthesis from text has been a highly active research area in the multimedia and computer vision community. This task is typically set up as a conditional image generation problem where a Generative Adversarial Network (GAN) [15] is learned to generate realistic looking images according to the text description in the format of natural languages [37, 39, 52, 71, 74, 84] or scene graphs [26, 42, 65, 72], etc.

In this paper, we study *how to manipulate image content through complex text instruction*. In this multimodal task, a user is able to apply various changes to a reference image by sending text instructions. For example, Figure 1 shows the generated images by the model for three types of instructions: 1) adding a new object at a location, 2) removing an object, and 3) changing the object's attributes (size, shape, color, etc). This concept was first raised in Schmandt and Hultheen's paper [61] and was extended to industrial applications such as PhotoShop through voice commands [59].

The task studied in this paper is inspired by cross-modal image retrieval – a cornerstone in many image retrieval tasks such as product search [6, 7, 16, 17, 32, 68, 80]. In this retrieval setting [6, 7, 68], users search an image database using a multimodal query that is

formed of an image plus some text that describes complex modifications to the input image. This retrieval problem is essentially the same as ours except we aim at generating as opposed to retrieving the target image. Notably, as will be shown in Section 4.2, the generated image can be used to as a query to retrieve the target images with competitive recall, thereby providing a more explainable search experience that allows users to inspect the search results before the retrieval.

The closest related problem to ours is text-guided image manipulation (e.g., [37, 52]). However, text instructions in existing works are limited in complexity and diversity as they mainly comprise descriptive attributes, lacking specific actions such as “add” or “remove” an object. In contrast, the considered text instructions in our paper cover three representative operations “add”, “modify”, and “remove” and involve adjectives (attributes), verbs (actions) and adverbs (locations) describing the intricate change to one of the objects in the reference image. Sequential image generation methods [4, 9, 12]) are also related. For example, GeNeVA [12] generates an image by adding objects to a blank canvas following the step-by-step instructions. Different from ours, these works tackle a different challenge, i.e., temporal modeling of the sequential image generation process.

The main challenge in our problem is how to model the *complex text instructions* for conditional image manipulation. To this end, we propose a simple yet highly effective approach called Text-Instructed Manipulation GAN or TIM-GAN. The key idea is to treat language as *neural operators* to locally modify the image feature for synthesizing the target image. The text neural operator decomposes the feature modification procedure into two stages: where and how to edit the image feature. For “where to edit”, we use attention mechanisms to ground words to a spatial region in the image. For “how to edit”, we introduce a text-adaptive network to generate different transformation for varying instructions. Since similar instructions perform similar operations, this design allows certain neurons to be shared among similar instructions, while still being able to distinguish among different operations.

We conduct extensive experiments on public datasets to demonstrate the three merits of the proposed method. First, it generates high-fidelity images, outperforming recent competitive baselines by a large margin. Second, the user studies confirm that the generated images are more semantically relevant to the target images. Third, the generated image, when used as the query for image-to-image retrieval, leads to not only promising retrieval recalls but also a more explainable search experience that allows users to inspect the results before the search. In addition, the ablation studies substantiate the performance gain stems from the proposed text operators. Code and models are released at <https://github.com/google/tim-gan>.

## 2 RELATED WORK

**Conditional generative adversarial networks.** Generative adversarial networks GANs [1, 2, 15, 50, 64] have made significant progress in recent years. Built on the basis of GANs, the *conditional* GAN aims to synthesize the image according to some input context. The input context can be images [22, 24, 34, 40, 51, 82], audio sequences [35], human poses [48], semantic segmentations [38, 54,

69], etc. Among them, text-to-image synthesis [4, 26, 36, 39, 41, 71–74, 84] learns a mapping from textual descriptions to images. Recently, GeNeVA [12] extended the mapping for iterative image generation in which new objects are added one-by-one to a blank canvas following textual descriptions. Different from text-to-image synthesis, the proposed problem takes multimodal inputs, aiming at learning to *manipulate* image content through text instructions.

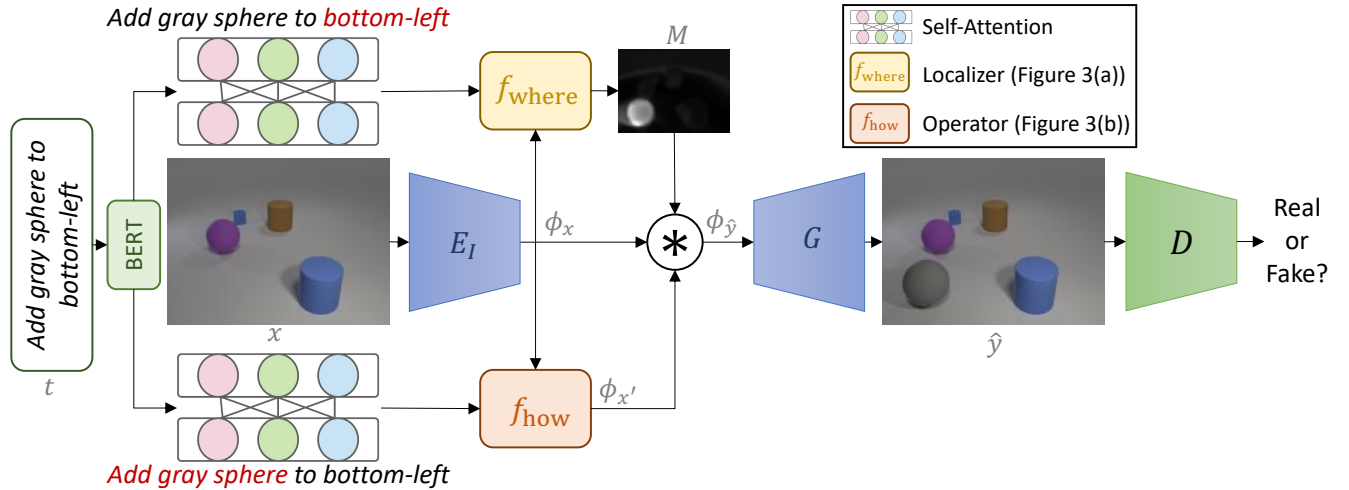
**Conditional image manipulation.** The research in this area aims to manipulate image content in a controlled manner. To enable user-guided manipulation, a variety of frameworks [3, 9, 21, 23, 33, 37, 43, 47, 52, 57, 62, 63, 70, 75, 77, 78, 81] have been proposed to study different control signals. For instance, Zhang et al. [78] and Zou et al. [86] used sparse dots and text respectively to guide the image colorization process. There are additional works on image manipulation by bounding boxes subsequently refined as semantic masks [20] or code [49]. Numerous image stylization [21, 43] and blending [23, 70] approaches augment the images by referencing an exemplar image. Other works include text-guided image inpainting [45, 76, 79] which use image caption to inpaint incomplete images. Closest to ours is the TA-GAN [52] scheme that takes the image caption as input to describe attributes for conditional image manipulation, followed by [37] and improved by [47]. In this work, we propose to manipulate the images according to the *complex* text instructions. Different from the image caption used by the TA-GAN, our instruction is more complex includes 1) three types of operations (“add”, “remove”, and “change”); 2) the explicit region information of the modification.

**Multimodal Feature Composition.** Another related area is multimodal feature composition which has been studied more extensively in other problems such as visual question answering [8, 30, 44, 53], visual reasoning [28, 60], image-to-image translation [34, 83], etc. Specifically, our method is related to feature-wise modulation, a technique to modulate the features of one source by referencing those from the other. Examples of recent contributions are: text image residual gating (TIRG) [68], feature-wise linear modulation (FiLM) [56], and feature-wise gating [14]. Among numerous works on multimodal feature composition, this paper compares the closely related methods including a strong feature composition method for image retrieval [68] and three competitive methods for conditional image generation [12, 52, 84].

## 3 METHODOLOGY

Our goal is to manipulate a given reference image according to the modification specified in the input text instruction from one of the three operations: “add”, “modify”, and “remove”. We approach this problem by modeling instructions as *neural operators* to modify the input image in the feature space. The text operator decomposes this process into two stages: where and how to edit the image feature. Thereafter, the edited feature is used to synthesize the target image by the generator of the GAN model.

An overview of the proposed TIM-GAN method is illustrated in Figure 2. Given the multimodal input: an image  $x$  and a text instruction  $t$ , our goal is to synthesize an image  $\hat{y}$  that is close to the ground-truth target image  $y$ . First, we extract the image feature  $\phi_x$  and the text features  $\phi_t$  where the text encoding comprises two heads producing  $\phi_t^{\text{where}}$  and  $\phi_t^{\text{how}}$  that encodes the *where* and *how* information about the text instruction, respectively. To indicate the



**Figure 2: Method overview.** Given an input image  $x$  and a text instruction  $t$ , the proposed TIM-GAN first predicts a spatial attention mask  $M$  (*where* to edit, Section 3.1) and a text operator  $f_{\text{how}}$  (*how* to edit, Section 3.2). The image feature  $\phi_x$  is then modified by the text operator  $f_{\text{how}}$  on the predicted mask  $M$ . Finally, the edited image  $\hat{y}$  is synthesized from the manipulated image feature  $\phi_{\hat{y}}$ .

region on the image  $x$  to be edited, we predict a spatial attention mask  $M$  from  $\phi_t^{\text{where}}$ . Then, we design a text-adaptive network that embodies a transformation ( $f_{\text{how}}$ ) for a text embedding ( $\phi_t^{\text{how}}$ ). Finally, using both the mask  $M$  and the embodied function  $f_{\text{how}}$ , the input image feature  $\phi_x$  is modified into  $\phi_{\hat{y}}$ , using which the resulting image  $\hat{y}$  is generated by  $G$ .

Formally, the image feature  $\phi_x$  is edited by the text operator by:

$$\phi_{\hat{y}} = \text{op}_{\text{text}}(\phi_x; t) \quad (1)$$

$$= (1 - M) \odot \phi_x + M \odot f_{\text{how}}(\phi_x, \phi_t^{\text{how}}; \Theta_{\text{how}}(t)), \quad (2)$$

where  $M = f_{\text{where}}(\phi_x, \phi_t^{\text{where}})$  is the learned spatial mask.  $\odot$  is element-wise dot product. The first term is a gated identity establishing the input image feature as a reference to the intended modified feature. Although the spatial attention or mask may not be a novel idea in image synthesis [5, 13, 46], we show that disentangling *how* and *where* in modification is essential for learning text operators that can be applied at various spatial locations. Our experimental results in Section 4.4 substantiate this claim.

The second term  $f_{\text{how}}$  embodies the specific computation to obtain the delta modification in the feature space. We introduce a text-adaptive network to execute different transformations for varying text inputs, where each text instruction is identified by a private set of parameters  $\Theta_{\text{how}}(t)$ , generated from  $\phi_t^{\text{how}}$ , and the remaining parameters are shared across all text instructions.

For training, we use the standard conditional GAN objective in the pix2pix [24] model, which consists of an adversarial loss  $\mathcal{L}_{\text{GAN}}$  and an  $\ell_1$  reconstruction loss called  $\mathcal{L}_{L1}$ . The weights to  $\mathcal{L}_{\text{GAN}}$  and  $\mathcal{L}_{L1}$  are set to 1 and 10, respectively. In the rest of this section, we will detail the computation of  $M$  and  $f_{\text{how}}$ .

### 3.1 Where to Edit: Spatial Mask

We use the scaled dot-product self-attention [67] to summarize the location-indicative, or locational words, in an instruction. Let  $S = [w_1, \dots, w_l] \in \mathbb{R}^{l \times d_0}$  denote the instruction where  $w_i \in \mathbb{R}^{d_0}$  is the word embedding [11] for the  $i$ -th word. The query, key and value in the attention are computed by:

$$Q = SW_Q, \quad K = SW_K, \quad V = SW_V \quad (3)$$

where  $W_Q, W_K, W_V \in \mathbb{R}^{d_0 \times d}$  are linear weight matrices to learn, and  $d$  is the output dimension. After reducing matrix  $Q$  to a column vector  $\hat{q}$  by average pooling along its first dimension, we obtain the attended text embedding by:

$$\phi_t^{\text{where}} = V^T \text{softmax}\left(\frac{K\hat{q}}{\sqrt{d}}\right), \quad (4)$$

in which the softmax function encourages higher attention weights over locational words. Likewise, we obtain the text feature  $\phi_t^{\text{how}}$  for salient operational words in the instruction (cf. Figure 6b), computed by a separate self-attention head.

We pass the image feature  $\phi_x$  to a convolution block (i.e., a ResBlock [18]) to get the output  $v \in \mathbb{R}^{H \times W \times C}$ . The spatial mask is then computed from  $\phi_t^{\text{where}}$  using image features as the context:

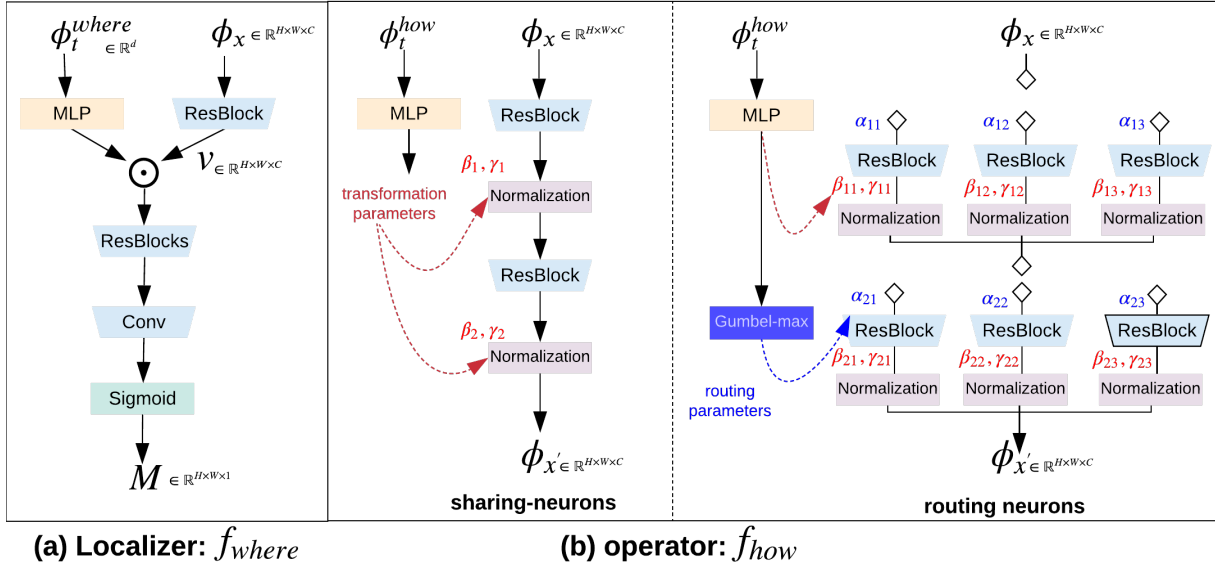
$$\begin{aligned} M &= f_{\text{where}}(\phi_x, \phi_t^{\text{where}}) \\ &= \sigma(W_m * (f_{\text{MLP}}(\phi_t^{\text{where}}) \odot v)) \in [0, 1]^{H \times W \times 1} \end{aligned} \quad (5)$$

where  $\sigma$  is the sigmoid function,  $*$  represents the 2d-convolution product with kernel  $W_m$  (cf. Figure 3a). We use two layers of the MLP with the ReLU activation.

During training, we compute an  $\ell_1$  loss to penalize the distance between the predicted mask  $M$  and the noisy true mask, and assign it the same weight as the  $\mathcal{L}_{L1}$  reconstruction loss. Note that computing this loss needs no additional supervision as the noisy mask is automatically computed by comparing the difference between the input and ground-truth training images.

### 3.2 How to Edit: Text-Adaptive Transformation

Text instructions are not independent. Similar instructions perform similar operations. For instance, “add a large cylinder” and “add a red cylinder” should perform virtually the same transformation except for the attribute part. Motivated by this idea, we design a text-adaptive network where each text instruction is instantiated by a few private parameters while the rest of the network parameters



**Figure 3: Where and how to edit. (a) The calculation of spatial mask  $M$  from text feature  $\phi_t^{where}$  and image feature  $\phi_x$ . (b) The proposed text-adaptive transformation network, parameterized by  $(\alpha, \beta, \gamma)$  generated from text feature  $\phi_t^{how}$ .**

are being shared across all text instructions. Below, we discuss two types of text-adaptive strategies.

**Sharing-Neurons.** In this strategy, every neuron in the network is shared among all the text instructions. An individual text is identified as a private set of parameters, i.e.,  $\Theta_{how}(t)$  in Equation 1, calculated from:

$$\Theta_{how}(t) = f_{MLP}(\phi_t^{where}) = \{(\beta_i, \gamma_i) | \gamma_i, \beta_i \in \mathbb{R}^p, i \in [1, l]\}, \quad (6)$$

where  $l$  is the total number layers of the text-adaptive network illustrated in Figure 3b. Each block consists of a conv layer followed by an instance normalization layer [66].  $p$  is the number of feature channels of each block. From the input feature  $\phi_t^{how}$ , an MLP layer is used to generate  $\beta$  and  $\gamma$  to perform text-specific batch normalization after the ResNet block. Our idea is partially inspired by the style transfer method [21].

**Routing-Neurons.** We find the above strategy works well in practice but is computationally expensive to scale up. We discuss an alternative strategy to apply text-adaptive transformation inside a routing network [58] where the text feature is used to dynamically select and execute a sequence of neural blocks (or a path). As a result, we call it routing-neurons strategy.

It is worth noting that our intention is not to compete the routing-neurons strategy with the sharing-neurons strategy because the former often does not lead to further performance gains. Yet, our goal is to show a scalable approach that efficiently increases the learning capacity of text operators, while still allowing certain neurons to be shared among similar instructions.

The text-adaptive network is shown in Figure 3b which has  $l$  layers of  $m$  blocks of identical structures. The routing parameter  $\alpha_i$  decides to connect or disconnect a block in a layer. A text instruction is hence parameterized by an additional series of  $\alpha$ :

$$\Theta_{how}(t) = \{(\alpha_i, \beta_i, \gamma_i) | \alpha_i \in [0, 1]^m, \gamma_i, \beta_i \in \mathbb{R}^{m \times p}, i \in [1, l]\}, \quad (7)$$

where  $\alpha_i, \beta_i, \gamma_i$  are all generated by the MLPs from  $\phi_t^{how}$ .

For efficiency, the path selector  $\alpha$  needs to take only discrete values. We employ the Gumbel-Softmax trick [25] to sample a block from a categorical distribution. Let  $\pi \in \mathbb{R}_{>0}^m$  be the categorical variable with probabilities  $P(\alpha = i) \propto \pi_i$ , i.e., the probability for selecting block  $i$ . We have:

$$\arg \max_i [P(\alpha = i)] = \arg \max_i [g_i + \log \pi_i] = \arg \max_i [\hat{\pi}_i], \quad (8)$$

where  $g_i = -\log(-\log(u_i))$  is a re-parameterization term, and  $u_i \sim \text{Uniform}(0, 1)$ . To make it differentiable, the softmax operation is used to compute  $\alpha = \text{softmax}(\hat{\pi}/\tau)$ , where the temperature  $\tau$  is set small to encourage  $\alpha$  being unimodal.

**Feature Modification.** Finally, the  $f_{how}$  function in the text operator (cf. Equation 2) is then calculated from:

$$f_{how}(\phi_x) = a^{(l+1)}, \quad (9)$$

$$a^{(i+1)} = \sum_{j=1}^m \alpha_{ij} (\gamma_{ij} \frac{o_{ij} - \mu(o_{ij})}{\delta(o_{ij})} + \beta_{ij}) \quad \forall i \in [1, l], \quad (10)$$

$$a^{(1)} = \phi_x, \quad (11)$$

where  $o_{ij}$  is the output of the  $j$ -th conv block in layer  $i$ .  $a^{(i)}$  is the activation of the  $i$ -th layer.  $\delta$  and  $\mu$  compute channel-wise mean and variance across spatial dimensions, and are applied at test time unchanged. Equation 9 details the feature modification step for both strategies. In particular, for the sharing-neurons strategy, we fix  $j$  to 1 and  $\alpha_{i1} = 1, \forall i \in [1, l]$  since there is only one path to choose from the network.

## 4 EXPERIMENTAL RESULTS

### 4.1 Setups

**Datasets. Clevr:** CSS dataset [68] was created for multimodal image retrieval using the Clevr toolkit [27]. The dataset contains 3-D synthesized images with multiple objects of varying colors, shapes, and sizes. Each training sample includes a reference image, a target image and a text specifying the modification from three types ‘‘add’’,

“remove”, “change” an object. The dataset includes 17K training and 17K test examples. **Abstract scene:** CoDraw [29] is a dataset built upon Abstract Scene [85] to illustrate a sequence of images of children playing in the park. For each sequence, there is a conversation between a Teller and a Drawer. The teller gives step-by-step instructions on how to add new content to the current image. Note its text is limited to the “add” operation. To adapt it to our problem, we extract the image and text of a single step. The dataset consists of 30K training and 8K test examples. **Cityscapes.** We create a new dataset of semantic segmentation from the Cityscapes dataset [10]. There are four types of text modifications: “add”, “remove”, “pull an object closer”, and “push an object away”. The ground-truth images are manually generated, according to the text instruction, by pasting desired objects onto the image at appropriate positions. The dataset consists of 20K training and 3K test examples.

Limited by suitable datasets, related works [12, 41] were only able to test on synthetic images (cf. more discussions in the supplementary material). In this paper, we extend our method to manipulate semantic segmentation in Cityscapes, and demonstrate the potential of our method for synthesizing RGB images from the modified segmentation mask.

**Baselines.** We compare with four baseline approaches. All methods are trained and tested on the same datasets, implemented using their official code or adapted official code. More details about the baseline comparison are discussed in the supplementary material.

- **DM-GAN:** The DM-GAN [84] model is a recent text-to-image synthesis framework. To adapt it to our task, we use our image encoder to extract the image feature and concatenate it with its original text feature as its input signal.
- **TIRG-GAN:** TIRG [68] is a competitive method for the cross-modal image retrieval task. It takes the same input as ours but only produces the image feature for retrieval. We build a baseline TIRG-GAN based on TIRG by using our image generator  $G$  to synthesize the image from the feature produced by the TIRG model.
- **TA-GAN:** TA-GAN [52] learns the mapping between the captions and images. The image manipulation is conducted by changing the text caption of the image. Since there is no image caption in our task, we concatenate the pre-trained features of the input image and text instruction as the input caption feature for the TA-GAN model.
- **GeNeVA:** GeNeVA [12] learns to generate the image step-by-step according to the text description. Its main focus is modeling the sequential image generation process. Nevertheless, to adapt it to take the same input as all the other methods, we use it for single-step generation over the real input image.

We select the above baseline methods because each of them represents the recent approach for the related problems of (a) text-to-image synthesis (DM-GAN), (b) multimodal retrieval (TIRG), (c) caption-based image manipulation (TA-GAN), and (d) sequential image generation (GeNeVA).

**Evaluation Metrics.** We employ two common metrics: Fréchet Inception Distance score (FID) [19] and retrieval recall. The former is used to measure the realism of the generated images, and the retrieval recall assesses the semantic relevance between the generated and the true target image.

To compute the retrieval recall, following [68, 71], we use the generated image as a query to retrieve the target images in the test set. For simplicity, we compute the cosine similarity between the features of the query and target images where the feature embeddings are obtained by an autoencoder pre-trained on each dataset.

**Implementation Details.** We implement our model in Pytorch [55]. For the image encoder  $E_I$ , we use three down-sampling convolutional layers followed by Instance Normalization with ReLU. We construct the generator  $G$  by using two residual blocks followed by three transposed-convolutional layers with Instance Normalization. For both  $E_I$  and  $G$ , we use 3x3 kernels and a stride of 2. For the text encoder  $E_t$ , we use the BERT [11] model. The encoded image has 256 feature channels and the attended text embedding dimension is  $d = 512$ . By default we use the routing-neurons strategy where the routing network has  $l = 2$  layers and  $m = 3$  blocks for each layer. The parameters in the image encoder  $E_i$  and decoder  $G$  are initialized by training an image autoencoder for 30 epochs. Then, we fix  $E_i$ 's parameters and optimize the other parts of the network in the end-to-end training for 60 epochs. For training, we use the Adam optimizer [31] with a batch size of 16, a learning rate of 0.002, and exponential rates of  $(\beta_1, \beta_2) = (0.5, 0.999)$ .

## 4.2 Main Results

The main results are shown in Table 1, where the Recall@ $N$  column indicates the recall ( $\times 100$ ) of the true target image in the top- $N$  retrieved images. The proposed method performs favorably against all baseline approaches across datasets. Although DM-GAN appears to generate more realistic images on the Clevr dataset, its retrieval scores are very poor ( $< 2\%$ ). This result indicates that it is deceiving to make comparisons only using FID because lower FIDs can be trivially obtained by merely copying the input image without any modifications.

Qualitative results are shown in Figure 4. As shown, TA-GAN and TIRG-GAN tend to copy the input images. DM-GAN often generates random objects following similar input layouts. GeNeVA can make local modifications to images, but often does not follow the text instructions. In contrast, our model generates images guided by the text instructions with greater fidelity and semantic relevance to the true target image.

We use the generated image by our model as a query to retrieve the target image. Figure 7 shows the top-5 returned images retrieved by our generated image on the Clevr dataset including two successful cases (the first 2 rows in Figure 7) and two failure cases. There is tangible resemblance between the generated query image and the true target image. This observation is consistent with the quantitative results presented in Table 1.

Figure 6 illustrates our intermediate results for where and how to edit, where the learned attention weights for the text and spatial mask are visualized. Generally, the attentions agree with our perception about the task as the self-attentions focus on locational and operational words in the text instruction, respectively, and the spatial attentions capture the intended area for modification.

The above quantitative results, in terms of both FID and retrieval score, substantiate that our method's efficacy in generating high-fidelity and semantically relevant images.



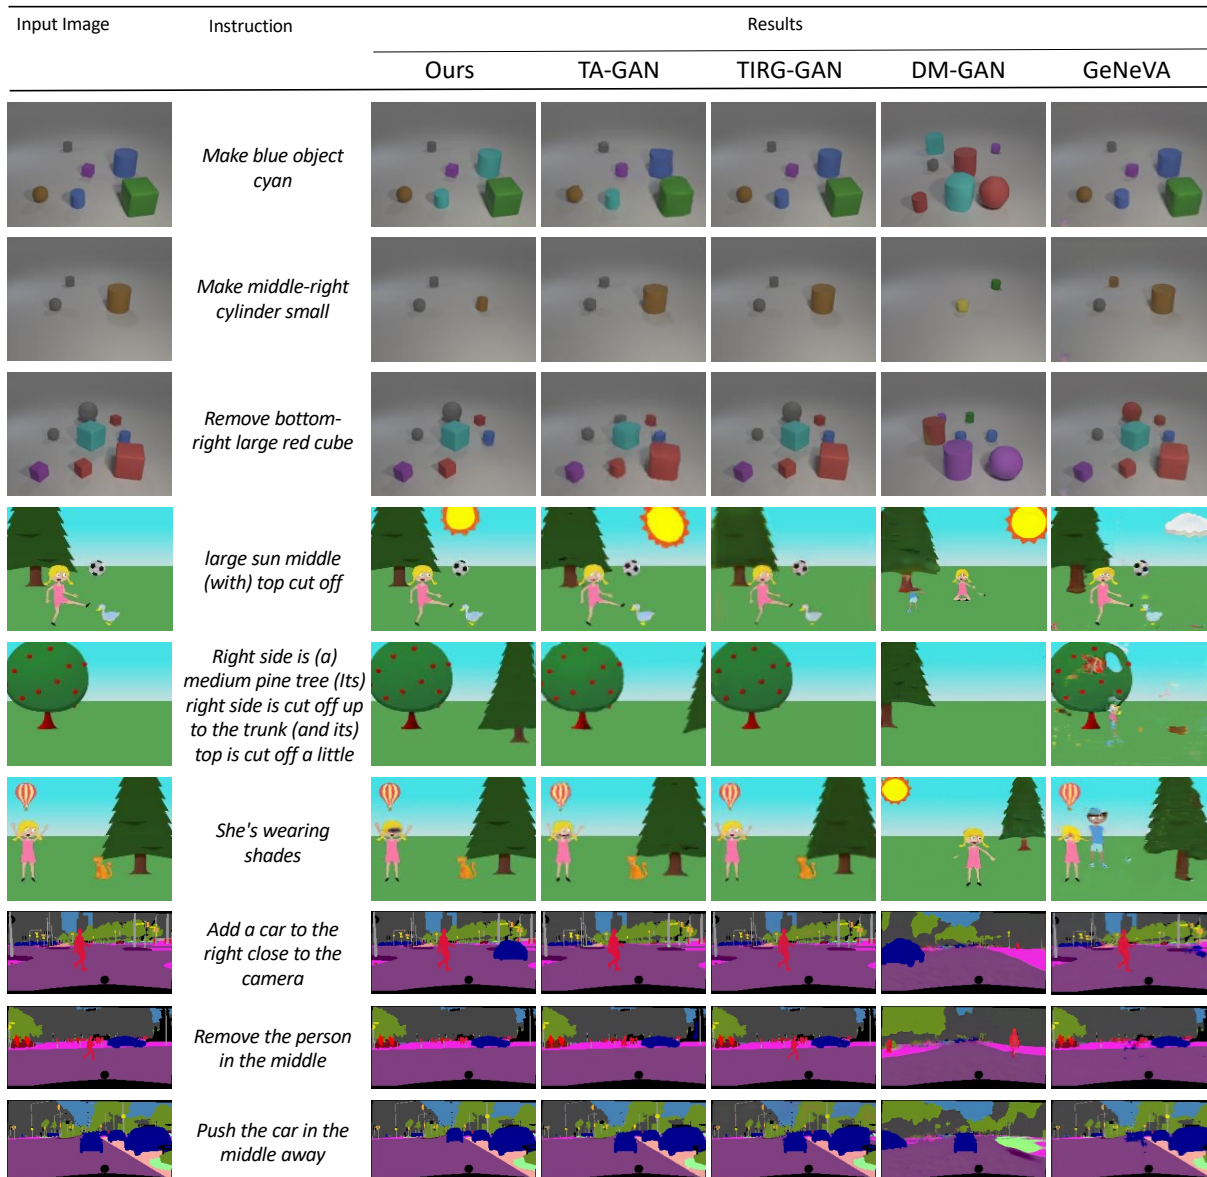


Figure 4: Selected generation results. We show the manipulation results by different approaches on the Clevr (top), Abstract scene (middle), and Cityscapes (bottom) datasets.

Table 1: Quantitative comparisons. We use the FID scores to measure the realism of the generated images, and the retrieval score (RS) to estimate the correspondence to text instructions.

Method	Clevr			Abstract scene			Cityscape		
	FID ↓	Recall@1 ↑	Recall@5 ↑	FID ↓	Recall@1 ↑	Recall@5 ↑	FID ↓	Recall@1 ↑	Recall@5 ↑
DM-GAN	<b>27.9</b>	1.6±0.1	5.6±0.1	53.8	2.1±0.1	6.6±0.1	18.7	4.6±0.2	15.7±0.2
TIRG-GAN	34.0	<u>48.5±0.2</u>	<u>68.2±0.1</u>	52.7	23.5±0.1	38.8±0.1	<u>6.1</u>	25.0±0.3	<u>88.9±0.3</u>
TA-GAN	58.8	40.8±0.1	64.1±0.1	<u>44.0</u>	<u>26.9±0.2</u>	<u>46.3±0.1</u>	6.7	<u>36.8±0.4</u>	79.8±0.3
GeNeVA	46.1	34.0±0.1	57.3±0.1	72.2	17.3±0.2	31.6±0.2	10.5	14.5±0.4	46.1±0.3
Ours	<u>33.0</u>	<b>95.9±0.1</b>	<b>97.8±0.1</b>	<b>35.1</b>	<b>35.4±0.2</b>	<b>58.7±0.1</b>	<b>5.9</b>	<b>77.2±0.4</b>	<b>99.9±0.1</b>
Real images	17.0	100	100	14.0	100	100	4.4	100	100




Figure 5: User preference studies. We present manipulated images on the Clevr and abstract scene datasets and ask the users to select the one which (a) is more realistic and (b) is more semantically relevant to the ground-truth image.

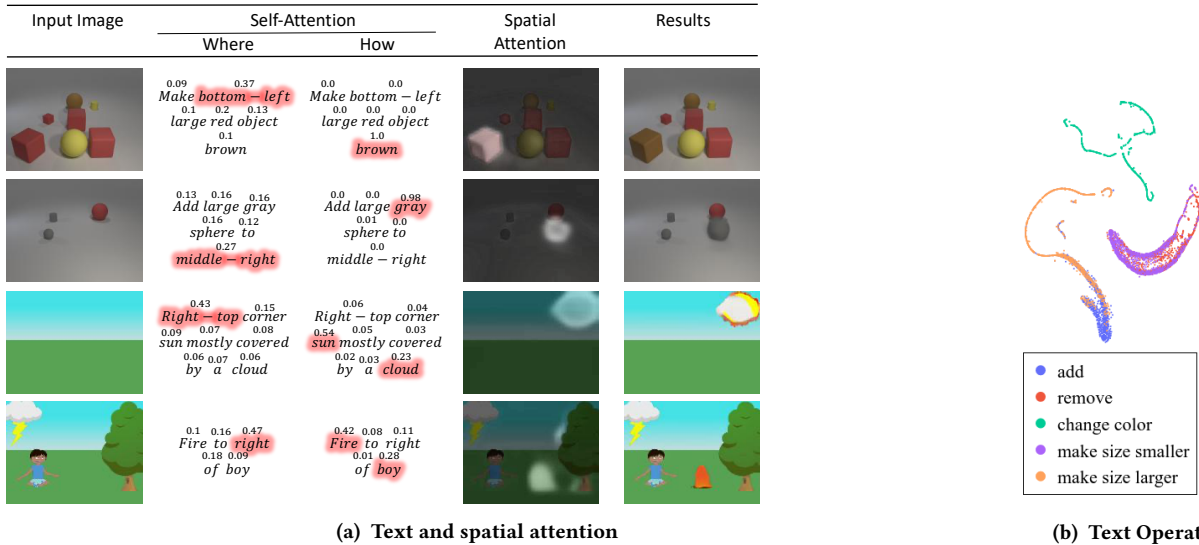


Figure 6: Where and how to edit. (a) Predicted self-attention weights and spatial attention masks. The self-attention weights are labeled above each word, and highlighted if the weights are greater than 0.2. (b) T-SNE visualization of the routing parameters  $\alpha$  for various types of text instructions on the Clevr dataset.

### 4.3 User preference study

We conduct two user studies to verify the visual quality and semantic relevance of the generated content. Given a pair of images generated by two compared methods, users are asked to choose (1) which one looks more realistic while ignoring the input image and text; (2) which one is more relevant to the text instruction by comparing the content of the generated and the ground-truth image. In total, we collected 960 answers from 30 users.

As shown in Figure 5, the proposed TIM-GAN outperforms other methods in both metrics with a statistically significant margin. The above results are consistent with the quantitative results in Table 1, which validate our method’s superior performance in generating not only realistic but also semantically relevant images.

### 4.4 Ablation studies

We test various ablations of our model to validate our design decisions by either leaving the module out from the full model or replacing it with an alternative module.

**Necessity of disentangling how and where to edit.** Our method is built upon a key idea to disentangle how and where to edit. To validate this design, we compare with two entangled text operators

in Table 2. The first removes the “where” information from the full model by replacing the spatial mask with an identity matrix. The second keeps the spatial mask but discards the “how” information by dropping the text-adaptive parameters from the  $f_{\text{how}}$  function. The inferior performance validates the necessity of disentangling how and where to edit. Note that replacing either module leads to worse performance than the baseline methods in Table 1, which indicates the performance gain is primarily from the proposed text operator as opposed to circumferential factors such as the network backbone or word embedding.

Table 2: Ablation on the disentangled text operator.

Method	Clevr		Abstract scene	
	FID ↓	Recall@1 ↑	FID ↓	Recall@1 ↑
Disentangled text operator	<b>33.0</b>	<b>95.9±0.1</b>	<b>35.1</b>	<b>35.4±0.2</b>
Entangled (no mask)	34.8	81.7±0.1	48.7	28.7±0.1
Entangled (no text-adaptive)	45.9	29.9±0.2	37.4	33.1±0.2

$f_{\text{how}}$  function. This experiment compares different designs of the  $f_{\text{how}}$  function. Three alternative models are considered from simple text-and-image feature concatenation, feature addition to the

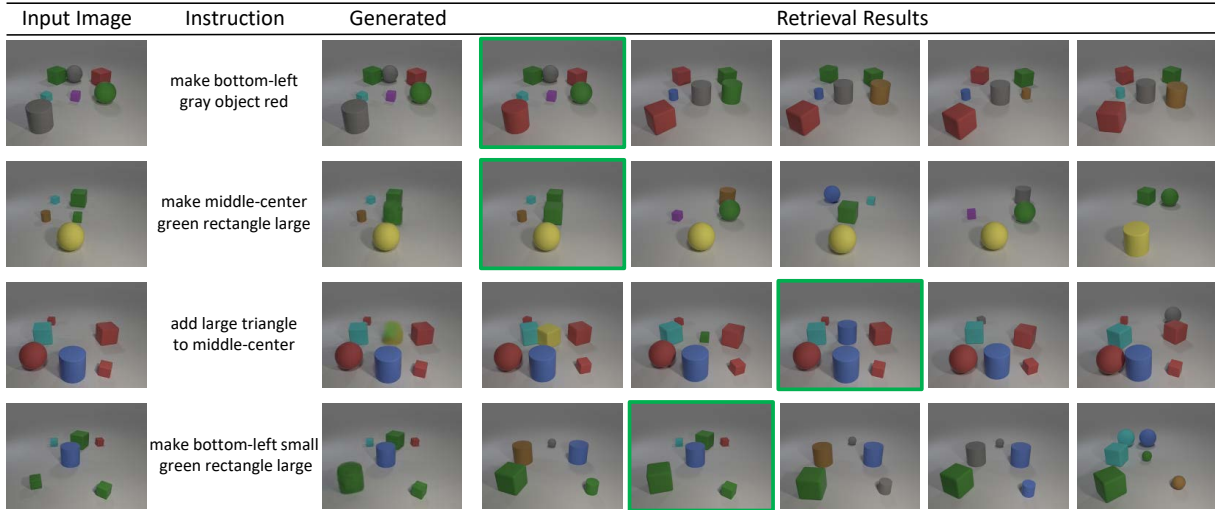


Figure 7: Top-5 images retrieved by our generated image (used as the query image). Third column shows the generated (fake) image by our model. Column 4-8 show the top-5 retrieved real images. The true target is highlighted in green.

more recent TIRG fusion [68]. All models use the same #layers, #parameters, and spatial mask. Figure 8 shows the results. The simple feature addition performs similarly on Clevr but is about 3% worse than our method on the Abstract Scene dataset.




Figure 8: Comparison between different  $f_{how}$  functions.

**With or without routing.** By default, we use the routing-neurons strategy. As discussed in Section 3.2, we do not intent to compete the routing and non-routing strategies because of their similar performances. Nevertheless, we fix the remaining network and study the sharing-neurons (without routing) strategy as well as another softmax routing strategy that computes the continuous routing parameter by the vanilla softmax function. All strategies use two layers and the routing strategies have 3 identical blocks per layer. Table 3 lists the results, where the FLOPs and Params column show the number of floating point operations and the number of network parameters of the text operator. We only count FLOPs and Params of the text operator since the remaining parts of the network are shared among strategies.

As discussed in Section 3.2, the routing-neurons strategy leads to no improvements but demonstrates two benefits. First, it allows for scaling up the network parameters at a marginal computational cost. Routing-neurons incurs similar #FLOPs as Sharing-neurons but inflates #Params by 3 fold. Second, the routing-neurons strategy enables neural blocks to be shared among similar text operators.

Figure 6b shows the t-SNE plot of its routing parameters  $\alpha$ . It is interesting to find this strategy automatically uncovers the subtle relationship between instructions. For instance, “add” and “make size larger” operators are closer and share more neural blocks.

Table 3: Ablation on sharing-neurons and routing-neurons strategies. Only the FLOPs (in Billion) and Params (in Million) of the text operator are compared.

Method	FLOPs	Params	Clevr	
			FID ↓	Recall@1 ↑
Sharing-Neurons	4.08B	4.46M	33.0	95.8±0.1
Routing-Neurons	4.09B	13.91M	33.0	95.9±0.1
Routing (Softmax)	12.24B	13.91M	33.0	95.4±0.1

## 5 CONCLUSION AND FUTURE WORK

In this paper, we studied a conditional image generation task that allows users to edit an input image using complex text instructions. We proposed an approach modeling text instructions as neural operators to locally modify the image feature. Our method decomposes “where” from “how” to apply the modification based on the design of text-adaptive networks. We evaluate our method on one real-world and two synthetic datasets, and obtain promising results with respect to three metrics on image quality, semantic relevance, and retrieval performance.

In the future, we plan to extend our work on more real-world datasets. Unfortunately, suitable evaluation benchmarks are currently unavailable for real-world RGB images. Therefore, one has to establish an evaluation benchmark of parallel triples of the reference RGB image, target RGB image, and text instruction. Following [47], we also hope to explore unsupervised training of our model on unpaired text and image data.



## REFERENCES

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein gan. In *International Conference on Machine Learning*.
- [2] Andrew Brock, Jeff Donahue, and Karen Simonyan. 2019. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*.
- [3] Huiwen Chang, Jingwan Lu, Fisher Yu, and Adam Finkelstein. 2018. Pairedcyclegan: Asymmetric style transfer for applying and removing makeup. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [4] Shizhe Chen, Bei Liu, Jianlong Fu, Ruihua Song, Qin Jin, Pingping Lin, Xiaoyu Qi, Chunting Wang, and Jin Zhou. 2019. Neural storyboard artist: Visualizing stories with coherent image sequences. In *ACM International conference on Multimedia*. 2236–2244.
- [5] Xinyuan Chen, Chang Xu, Xiaokang Yang, and Dacheng Tao. 2018. Attention-gan for object transfiguration in wild images. In *European Conference on Computer Vision*. 164–180.
- [6] Yanbei Chen and Loris Bazzani. 2020. Learning Joint Visual Semantic Matching Embeddings for Language-guided Retrieval. In *European Conference on Computer Vision*.
- [7] Yanbei Chen, Shaogang Gong, and Loris Bazzani. 2020. Image search with text feedback by visiolinguistic attention learning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [8] Yanbei Chen, Shaogang Gong, and Loris Bazzani. 2020. Image Search with Text Feedback by Visiolinguistic Attention Learning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [9] Yu Cheng, Zhe Gan, Yitong Li, Jingjing Liu, and Jianfeng Gao. 2020. Sequential attention GAN for interactive image editing. In *Proceedings of the 28th ACM International Conference on Multimedia*. 4383–4391.
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. 2016. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [12] Alaaeldin El-Nouby, Shikhar Sharma, Hannes Schulz, Devon Hjelm, Layla El Asri, Samira Ebrahimi Kahou, Yoshua Bengio, and Graham W. Taylor. 2019. Tell, Draw, and Repeat: Generating and modifying images based on continual linguistic instruction. In *IEEE International Conference on Computer Vision*.
- [13] Hajar Emami, Majid Moradi Aliabadi, Ming Dong, and Ratna Babu Chinnam. 2020. Spa-gan: Spatial attention gan for image-to-image translation. *IEEE Transactions on Multimedia* 23 (2020), 391–401.
- [14] Arnab Ghosh, Richard Zhang, Puneet K Dokania, Oliver Wang, Alexei A Efros, Philip HS Torr, and Eli Shechtman. 2019. Interactive sketch & fill: Multiclass sketch-to-image translation. In *IEEE International Conference on Computer Vision*.
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Neural Information Processing Systems*.
- [16] Xiaoxiao Guo, Hui Wu, Yu Cheng, Steven Rennie, Gerald Tesaro, and Rogerio Feris. 2018. Dialog-based interactive image retrieval. In *Neural Information Processing Systems*.
- [17] Xiaoxiao Guo, Hui Wu, Yupeng Gao, Steven Rennie, and Rogerio Feris. 2019. Fashion IQ: A New Dataset towards Retrieving Images by Natural Language Feedback. *arXiv preprint arXiv:1905.12794* (2019).
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [19] 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 *Neural Information Processing Systems*.
- [20] Seunghoon Hong, Xinchen Yan, Thomas S Huang, and Honglak Lee. 2018. Learning hierarchical semantic image manipulation through structured representations. In *Neural Information Processing Systems*.
- [21] Xun Huang and Serge Belongie. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE International Conference on Computer Vision*.
- [22] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. 2018. Multimodal unsupervised image-to-image translation. In *European Conference on Computer Vision*.
- [23] Wei-Chih Hung, Jianming Zhang, Xiaohui Shen, Zhe Lin, Joon-Young Lee, and Ming-Hsuan Yang. 2018. Learning to Blend Photos. In *European Conference on Computer Vision*.
- [24] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [25] Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*.
- [26] Justin Johnson, Agrim Gupta, and Li Fei-Fei. 2018. Image Generation from Scene Graphs. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [27] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [28] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Judy Hoffman, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Inferring and executing programs for visual reasoning. In *IEEE International Conference on Computer Vision*.
- [29] Jin-Hwa Kim, Nikita Kitaev, Xinlei Chen, Marcus Rohrbach, Byoung-Tak Zhang, Yuandong Tian, Dhruv Batra, and Devi Parikh. 2017. CoDraw: Collaborative drawing as a testbed for grounded goal-driven communication. *arXiv preprint arXiv:1712.05558* (2017).
- [30] Jin-Hwa Kim, Sang-Woo Lee, Donghyun Kwak, Min-Oh Heo, Jeonghee Kim, Jung-Woo Ha, and Byoung-Tak Zhang. 2016. Multimodal Residual Learning for Visual QA. In *Neural Information Processing Systems*.
- [31] Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- [32] Adriana Kovashka, Devi Parikh, and Kristen Grauman. 2012. Whittlesearch: Image search with relative attribute feedback. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [33] Gierad P Laput, Mira Dontcheva, Gregg Wilensky, Walter Chang, Aseem Agarwala, Jason Linder, and Eytan Adar. 2013. PixelTone: A multimodal interface for image editing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2185–2194.
- [34] Hsin-Ying Lee, Hung-Yu Tseng, Qi Mao, Jia-Bin Huang, Yu-Ding Lu, Maneesh Singh, and Ming-Hsuan Yang. 2020. DriT++: Diverse image-to-image translation via disentangled representations. *International Journal of Computer Vision* (2020).
- [35] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. 2019. Dancing to Music. In *Neural Information Processing Systems*.
- [36] Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip Torr. 2019. Controllable text-to-image generation. In *Neural Information Processing Systems*.
- [37] Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip HS Torr. 2020. ManiGAN: Text-Guided Image Manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [38] Ke Li, Tianhao Zhang, and Jitendra Malik. 2019. Diverse image synthesis from semantic layouts via conditional IMLE. In *IEEE International Conference on Computer Vision*.
- [39] 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 *IEEE Conference on Computer Vision and Pattern Recognition*.
- [40] Yuhang Li, Xuejin Chen, Binxin Yang, Zihan Chen, Zhihua Cheng, and Zheng-Jun Zha. 2020. DeepFacePencil: Creating Face Images from Freehand Sketches. In *Proceedings of the 28th ACM International Conference on Multimedia*. 991–999.
- [41] Yitong Li, Zhe Gan, Yelong Shen, Jingjing Liu, Yu Cheng, Yuexin Wu, Lawrence Carin, David Carlson, and Jianfeng Gao. 2019. Storygan: A sequential conditional gan for story visualization. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [42] Yijun Li, Lu Jiang, and Ming-Hsuan Yang. 2021. Controllable and Progressive Image Extrapolation. In *IEEE/CVF Winter Conference on Applications of Computer Vision*.
- [43] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. 2018. A closed-form solution to photorealistic image stylization. In *European Conference on Computer Vision*.
- [44] Junwei Liang, Lu Jiang, Liangliang Cao, Yannis Kalantidis, Li-Jia Li, and Alexander G Hauptmann. 2019. Focal visual-text attention for memex question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 8 (2019), 1893–1908.
- [45] Qing Lin, Bo Yan, Jichun Li, and Weimin Tan. 2020. MMFL: Multimodal Fusion Learning for Text-Guided Image Inpainting. In *Proceedings of the 28th ACM International Conference on Multimedia*. 1094–1102.
- [46] Yu Lin, Yigong Wang, Yifan Li, Yang Gao, Zhuoyi Wang, and Latifur Khan. 2021. Attention-Based Spatial Guidance for Image-to-Image Translation. In *IEEE/CVF Winter Conference on Applications of Computer Vision*. 816–825.
- [47] Yahui Liu, Marco De Nadai, Deng Cai, Huayang Li, Xavier Alameda-Pineda, Nicu Sebe, and Bruno Lepri. 2020. Describe What to Change: A Text-guided Unsupervised Image-to-Image Translation Approach. In *Proceedings of the 28th ACM International Conference on Multimedia*. 1357–1365.
- [48] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. 2017. Pose guided person image generation. In *Neural Information Processing Systems*.
- [49] Jiayuan Mao, Xiuming Zhang, Yikai Li, William T Freeman, Joshua B Tenenbaum, and Jiajun Wu. 2019. Program-Guided Image Manipulators. In *IEEE International Conference on Computer Vision*.
- [50] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. 2017. Least squares generative adversarial networks. In *IEEE International Conference on Computer Vision*.

- [51] Youssef Alami Mejjati, Christian Richardt, James Tompkin, Darren Cosker, and Kwang In Kim. 2018. Unsupervised attention-guided image-to-image translation. In *Neural Information Processing Systems*.
- [52] Seonghyeon Nam, Yunji Kim, and Seon Joo Kim. 2018. Text-adaptive generative adversarial networks: manipulating images with natural language. In *Neural Information Processing Systems*.
- [53] Hyeonwoo Noh, Paul Hongsuck Seo, and Bohyung Han. 2016. Image question answering using convolutional neural network with dynamic parameter prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [54] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. 2019. Semantic image synthesis with spatially-adaptive normalization. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [55] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *NeurIPS workshop*.
- [56] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. Film: Visual reasoning with a general conditioning layer. In *AAAI Conference on Artificial Intelligence*.
- [57] Tiziano Portenier, Qiyang Hu, Attila Szabo, Siavash Arjomand Bigdeli, Paolo Favaro, and Matthias Zwicker. 2018. Faceshop: Deep sketch-based face image editing. *ACM Transactions on Graphics* 37, 4 (2018), 99.
- [58] Clemens Rosenbaum, Tim Klinger, and Matthew Riemer. 2018. Routing networks: Adaptive selection of non-linear functions for multi-task learning. In *International Conference on Learning Representations*.
- [59] Cella Lao Rousseau. 2017. Bye-bye basic editing, hello voice-controlled Photoshop! <https://www.imore.com/bye-bye-basic-editing-hello-voice-controlled-photoshop>
- [60] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. 2017. A simple neural network module for relational reasoning. In *Neural Information Processing Systems*.
- [61] Christopher Schmandt and Eric A Hulteen. 1982. The intelligent voice-interactive interface. In *Proceedings of the 1982 conference on Human factors in computing systems*. 363–366.
- [62] Seitaro Shinagawa, Koichiro Yoshino, Seyed Hossein Alavi, Kallirroi Georgila, David Traum, Sakriani Sakti, and Satoshi Nakamura. 2020. An Interactive Image Editing System Using an Uncertainty-Based Confirmation Strategy. *IEEE Access* 8 (2020), 98471–98480.
- [63] Hung-Yu Tseng, Matthew Fisher, Jingwan Lu, Yijun Li, Vladimir Kim, and Ming-Hsuan Yang. 2020. Modeling artistic workflows for image generation and editing. In *European Conference on Computer Vision*.
- [64] Hung-Yu Tseng, Lu Jiang, Ce Liu, Ming-Hsuan Yang, and Weilong Yang. 2021. Regularizing Generative Adversarial Networks under Limited Data. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [65] Hung-Yu Tseng, Hsin-Ying Lee, Lu Jiang, Ming-Hsuan Yang, and Weilong Yang. 2020. RetrieveGAN: Image synthesis via differentiable patch retrieval. In *European Conference on Computer Vision*.
- [66] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. 2017. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [67] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Neural Information Processing Systems*.
- [68] Nam Vo, Lu Jiang, Chen Sun, Kevin Murphy, Li-Jia Li, Li Fei-Fei, and James Hays. 2019. Composing text and image for image retrieval-an empirical odyssey. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [69] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2018. High-resolution image synthesis and semantic manipulation with conditional GANs. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [70] Huikai Wu, Shuai Zheng, Junge Zhang, and Kaiqi Huang. 2019. Gp-gan: Towards realistic high-resolution image blending. In *Proceedings of the 27th ACM international conference on multimedia*. 2487–2495.
- [71] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. 2018. AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [72] Li Yikang, Tao Ma, Yeqi Bai, Nan Duan, Sining Wei, and Xiaogang Wang. 2019. Pastegan: A semi-parametric method to generate image from scene graph. In *Neural Information Processing Systems*.
- [73] Mingkuan Yuan and Yuxin Peng. 2018. Text-to-image synthesis via symmetrical distillation networks. In *Proceedings of the 26th ACM international conference on Multimedia*. 1407–1415.
- [74] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. 2018. StackGAN++: Realistic image synthesis with stacked generative adversarial networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 8 (2018), 1947–1962.
- [75] Jichao Zhang, Yezhi Shu, Songhua Xu, Gongze Cao, Fan Zhong, Meng Liu, and Xueying Qin. 2018. Sparsely grouped multi-task generative adversarial networks for facial attribute manipulation. In *Proceedings of the 26th ACM international conference on Multimedia*. 392–401.
- [76] Lisai Zhang, Qingcai Chen, Baotian Hu, and Shuoran Jiang. 2020. Text-Guided Neural Image Inpainting. In *Proceedings of the 28th ACM International Conference on Multimedia*. 1302–1310.
- [77] Richard Zhang, Phillip Isola, and Alexei A Efros. 2016. Colorful Image Colorization. In *European Conference on Computer Vision*.
- [78] Richard Zhang, Jun-Yan Zhu, Phillip Isola, Xinyang Geng, Angela S Lin, Tianhe Yu, and Alexei A Efros. 2017. Real-time user-guided image colorization with learned deep priors. *ACM Transactions on Graphics* 36, 4 (2017), 1–11.
- [79] Zijian Zhang, Zhou Zhao, Zhu Zhang, Baoxing Huai, and Jing Yuan. 2020. Text-Guided Image Inpainting. In *Proceedings of the 28th ACM International Conference on Multimedia*. 4079–4087.
- [80] Bo Zhao, Jiashi Feng, Xiao Wu, and Shuicheng Yan. 2017. Memory-Augmented Attribute Manipulation Networks for Interactive Fashion Search. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [81] Na Zheng, Xuemeng Song, Zhaozheng Chen, Linmei Hu, Da Cao, and Liqiang Nie. 2019. Virtually trying on new clothing with arbitrary poses. In *Proceedings of the 27th ACM International Conference on Multimedia*. 266–274.
- [82] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision*.
- [83] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. 2017. Toward multimodal image-to-image translation. In *Neural Information Processing Systems*.
- [84] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. 2019. Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [85] C Lawrence Zitnick and Devi Parikh. 2013. Bringing semantics into focus using visual abstraction. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [86] Changqing Zou, Haoran Mo, Chengying Gao, Ruofei Du, and Hongbo Fu. 2019. Language-based colorization of scene sketches. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–16.