

# Learning to Globally Edit Images with Textual Description

Hai Wang<sup>†</sup> Jason D. Williams<sup>‡</sup> Sing Bing Kang<sup>§</sup>

**Abstract.** We show how we can globally edit images using textual instructions: given a source image and a textual instruction for the edit, generate a new image transformed under this instruction. To tackle this novel problem, we develop three different trainable models based on RNN and Generative Adversarial Network (GAN). The models (bucket, filter bank, and end-to-end) differ in how much expert knowledge is encoded, with the most general version being purely end-to-end. To train these systems, we use Amazon Mechanical Turk to collect textual descriptions for around 2000 image pairs sampled from several datasets. Experimental results evaluated on our dataset validate our approaches. In addition, given that the filter bank model is a good compromise between generality and performance, we investigate it further by replacing RNN with Graph RNN, and show that Graph RNN improves performance. To the best of our knowledge, this is the first computational photography work on global image editing that is purely based on free-form textual instructions.

## 1 Introduction

Consumers are increasingly relying on portable embedded devices such as smartphones and tablets for their everyday activities. These devices tend to have small form factors that preclude fine-grain spatial control using the display. Adding voice-based instruction (systems such as Siri, Cortana, and Alexa) significantly enhances the capabilities of such devices. An application that would significantly benefit is photo editing.

With few exceptions, interactive photo editing systems are primarily manual and often require significant display real estate for the controls. To allow a photo editing system to be voice-controlled, the mapping of voice to text to invocation of image operations requires domain-specific conversion of text to APIs. One solution is to handcraft this conversion by manually defining rules for editing effects (as was done in [1]). However, this approach is hard to scale.

---

<sup>†</sup> TTIC, Chicago, IL, 60637. USA. Email: haiwang@ttic.edu, work done at MSR

<sup>‡</sup> Apple, Cupertino, CA, 95014. USA, Email: jdw@alumni.princeton.edu, work done at MSR

<sup>§</sup> Microsoft Research, Redmond, WA, 98052, USA. Email: sbkang@microsoft.com

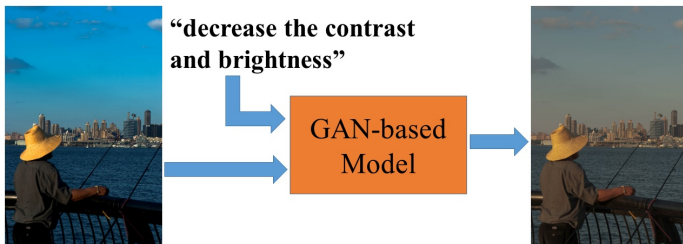


Fig. 1: Overview of our system. The inputs are an image and textual command, with the output being the result of applying the command to the input image.

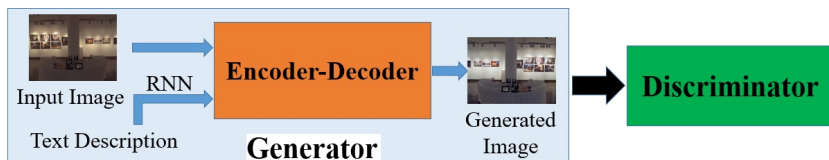


Fig. 2: Our GAN-based system.

In this paper, we demonstrate global image editing through text, as illustrated in Figure 1. Compared to other work [2,1], our system is end-to-end trainable and easier to extend, since it does not require significant handcrafting of rules. We designed three different models based on Generative Adversarial Network (GAN) [3]. Our main contributions are:

- We believe our work is the first to tackle the general image editing problem under free-form text descriptions.
- We collected a database of image transformation pairs and their corresponding textual descriptions.
- We designed three different models: handcrafted bucket-based model, pure end-to-end model, and filter bank based model. Experimental results demonstrate the effectiveness of our approaches.
- We are the first method to apply graph RNN to text-image synthesis and demonstrate its effectiveness.

In our work, we limit image editing to global transforms<sup>1</sup>.

## 2 Related work

In this section, we briefly describe two voice-assisted systems, namely, PixelTone [1] for image editing, and Image Spirit [2] for refining a parsed image. We also survey

<sup>1</sup> Code is available at [https://github.com/sohuren/Img\\_edit\\_with\\_text](https://github.com/sohuren/Img_edit_with_text).



representative approaches for automatic image editing (specifically, image enhancement and style transfer), joint image-language analysis, and techniques that use attention or graph RNN.

***PixelTone and Image Spirit.*** From application side, PixelTone [1] is the system most related to ours. It allows the user to edit the image through the voice command such as “change the t-shirt to blue”, after the t-shirt region is tagged. The system contains a speech recognition engine, a text analysis module, and an execution module. After converting the user’s voice command to text, the text analysis module produces the atomic operation which can be run by the execution module. The text analysis module is based on predefined rules and NLP techniques such as tokenization and part of speech tagging. One limitation is that the predefined rules are manually constructed.

Image Spirit [2] is a system that parses an image into regions with semantic labels, and allows the user to verbally refine the result. Typical verbal commands include correcting an object label and refining a specific label. Based on the initial image parsing result, Image Spirit updates the local relationship between different objects in an MRF [4] in response to the utterance input, resulting in the enhanced result. As with PixelTone, the commands are also predefined, and the scenario of refining the image parsing result is different from our image editing scenario.

Unlike PixelTone and Image Spirit, our approach does not rely on pre-define commands or rules, rather, it learns an end-to-end model which takes arbitrary text and learns corresponding transformations, based on a corpus.

***Image Manipulation with Language.*** Concurrent with our work, there are several techniques that address end-to-end trainable model for image manipulation with language [5,6,7]. Chen et al. [5] developed attentive models capable of combining text and image to produce a new image. To extract meaningful information from text, they use the attention mechanism [8]. They demonstrate two editing tasks of image segmentation and colorization with natural language; different losses are used for training the different tasks. The work of Seitaro et al. [6] is similar to Chen et al. [5], but they only consider MNIST dataset with instructions related to position moving such as “moving 6 to the bottom.” By creating the artificial dataset, they explored what the model can learn; as a side effect, training on an artificial dataset limits practicality.

By comparison, our model focuses on general textual instructions, which makes it extensible to different kinds of instructions. Further, instead of using attention mechanism [8,5], we use graph RNN [9]. Finally, our model is trained on real dataset collected from Amazon Mechanical Turk.

***Automatic Image Editing: Enhancement and Style Transfer.*** There are a number of approaches for automatic image enhancement. Machine learning techniques have been used to train on original-enhanced image pair databases for enhancing images [10,11,12,13]. The approach of [14] is based on a trained deep neural network to predict the enhanced image. Another form of image editing is style transfer (exemplified by [15,16,17,18]). Here, given one image and reference image, the goal is to generate the new image according to the reference image’s style. The mapping is purely image-based. Unlike our work, all these techniques do not act on a textual description.

**Joint Image and Language Analysis.** A significant amount of work has been done on joint image-language analysis. Topics in this space include image caption generation [19,20,21,22,23], video story telling [24,25], visual question answering [26,27,28,29], image retrieval under natural language [30], object retrieval under language [31,32,7], image synthesis from text [33,34,35,36] and referring expression generation [37,38,39].

The topics of image retrieval under natural language, object retrieval, image synthesis from text, and referring expression generation are most relevant to our work. Ulyanov et al. [30] use natural language to guide image retrieval; image-text correspondence is used to find a common embedding space. There are techniques that, given text and image, localize a target object as a bounding box [31,32] or segment [7] within the image. The work of Mirza and Osindero [37] addresses the problem of referring expression generation, i.e., given image and a bounding box, generate the expression that can describe it. The techniques of [39] and [38] generalize this problem in the context of reinforcement learning.

Given image objects and text, the approaches of [40] and [41] find the alignment between them. Kong et al. [42] find the alignment between text and RGB-D image, and use the text description to guide 3D semantic parsing. They show that image information helps to improve the language analysis result.

In [33,34], the output image is synthesized from noise vector and text description, but in our work, we begin from the original image and try to transform the image under text description. The technique of [33] generates a fixed size image while our output image size depends on the input size, which complicates the image generation problem. Additionally, we rely on basic image concepts such as saturation and brightness while the techniques of [33,34,35] analyzes image content (with only image concept involved being color).

In summary, all these techniques that focus on joint image-language analysis are not designed to transform the image using text. However, if we were to transform the image locally, as in “change the color of the dog on right to white”, we would need to first localize the dog before applying the color change. Here, techniques such as [31,32,7] would be good candidate components to add to our system.

**Attention and graph RNN.** Attention has been used in various joint image and text problems, and generally there are two different attention mechanism: attention between different tokens in text [5], and attention between tokens in text and pixels in image [22,35].

Graph RNN is first used for cross-sentence  $N$ -ary relation extraction [9], and it subsumes plain&tree RNN [43]. Briefly, a graph RNN generalizes a linear-chain RNN by incorporating arbitrary long-ranged dependencies besides word adjacency. A word might have precedents other than the prior word, and its LSTM unit is expanded to include one forget gate for each precedent. For efficient training, a graph is decomposed into a forward pass and a backward pass, each consisting of edges pointing forward and backward, respectively. Backpropagation is then conducted on these two directed acyclic graphs, similarly to BiLSTM. (If a graph LSTM contains no edges other than word adjacency, it reduces to BiLSTM.) Additional dependencies include syntactic dependencies, discourse relations, coreference, and connections between roots of adjacent sentences [9].

In our work, we handle only global editing; given our limited training data, how to extract meaningful semantics from text is crucial. As such, graph RNN might be a more natural choice than attention since they can utilize graph structure provided by parser [44]. To the best of our knowledge, this is the first time that graph structured RNN is used in a joint text and image analysis problem.

### 3 Model

Our goal is to use text and an image as input and generate a new image globally transformed under the text description. This problem is well-suited to the adversarial framework provided by Generative Adversarial Network [3,45]. The GAN objective function is a min-max problem, which is typically optimized in an alternating manner:

$$\min_{\theta_G} \max_{\theta_D} (\mathbb{E}_{x \sim p_{data}(x)} [\log D_{\theta_D}(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D_{\theta_D}(G_{\theta_G}(z)))]), \quad (1)$$

where  $D_{\theta_D}$  is the discriminator with parameter  $\theta_D$  and  $G_{\theta_G}$  is the generator with parameter  $\theta_G$ . The generator tries to confuse the discriminator while the discriminator differentiates between samples from true data distribution  $p_{data}$  and samples from the generator given noise data distribution  $p_z(z)$ .

In our work, the image transformation is achieved by the generator, which consists of an encoder-decoder architecture and a Recurrent Neural Network (RNN). As for the discriminator,  $p_z(z)$  is the original image while  $p_{data}$  is the corresponding edited image. The system is depicted in Figure 2.

We design three models, and each model handles the text information differently:

1. Hand-crafted bucket-based model, where similar image transformations are grouped prior to training as buckets. Each bucket has its own encoder-decoder architecture.
2. End-to-end model, with a single encoder-decoder architecture to handle the image and an RNN to handle text.
3. Filter-bank model, where transformations are specified as trained convolution filters.

All these models have exactly the same discriminator, and they only differ in the generator.

#### 3.1 Discriminator

We describe the discriminator first since it is same for all three models. We consider the conditional GAN (c-GAN) [45] where the loss is also conditioned on the input image. However, compared with [45], our discriminator also need to be text aware because the image is enhanced under the corresponding text description.

Our discriminator uses four inputs: original image  $I_{input}$ , ground truth image  $I_{gt}$  or generated image  $I_g$ , and the corresponding text  $T_{des}$ . As with [33], for each image pair, we also consider sampling random texts (see Section B in supplementary) to make the discriminator more text-aware.

Let  $h(x)$  be an encoding function (e.g., an RNN) which can encode text  $x$  into a vector. We first encode the text and down-sample the image before we depth-concatenate  $I_{input}$ ,  $I_g$ , and  $h(T_{des})$ , then we feed the resulting vector to the discriminator with a negative label. In contrast, we feed the triple  $I_{input}$ ,  $I_{gt}$  and  $h(T_{des})$  to the discriminator with a positive label.

Additionally, the triple  $I_{input}$ ,  $I_{gt}$  and  $h(T_{random})$  and  $I_{input}$ ,  $I_g$  and  $h(T_{random})$  are treated as negative instances. The discriminator loss is just summed over all instances.

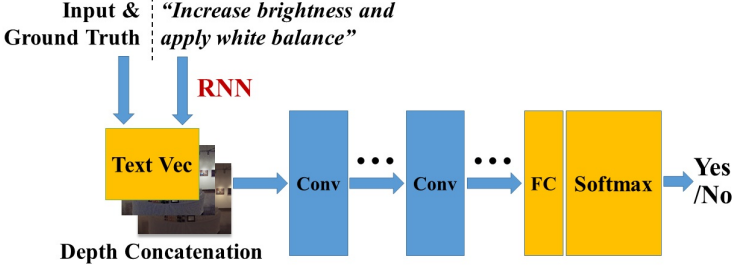


Fig. 3: Our discriminator architecture.

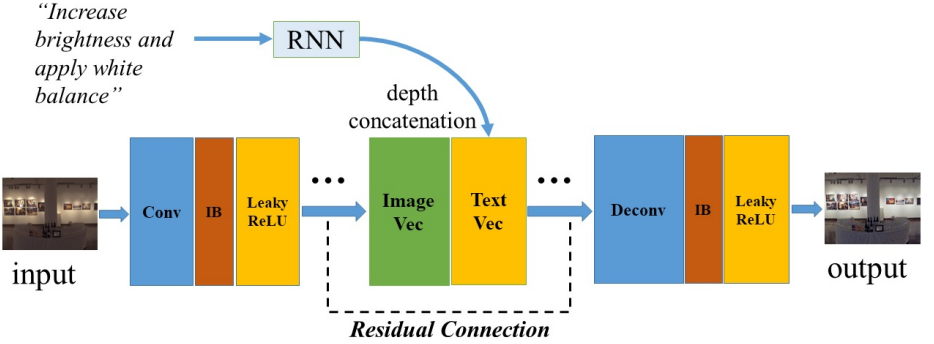


Fig. 4: Our end-to-end model.

The discriminator architecture is depicted in Figure 3. It has fewer layers, and each layer contains convolution, instance normalization [46], and activation function. Compared with [33], we use the sampled text in the context of c-GAN while [33] uses the sampled text in basic GAN. More details on the discriminator are provided in Section B of the supplementary file.

### 3.2 Generators

All three generators take an image and text as input and generate a corresponding transformed image. Before describing our three models, we define the loss function. Given original image  $I_{input}$ , generated image  $I_g$  and ground truth  $I_{gt}$ , we use the following losses to train the generator:

– **Content loss:**

$$l_{\text{content}} = \frac{1}{WHC} \sum_{c=1}^C |I_g^c - I_{gt}^c|_1, \quad (2)$$

where  $W$ ,  $H$ , and  $C$  are the image width, image height, and channel number, respectively.  $l_{\text{content}}$  measures  $l_1$  loss of the generated and ground truth images.

– **Adversarial loss:**

$$l_{\text{adversarial}} = 1 - \log(D_{\theta_D}(I_{input}, I_g, h(T_{des}))), \quad (3)$$

where  $h(T_{des})$  is defined in (6).  $l_{\text{adversarial}}$  comes from the discriminator; it measures the similarity of the generated image with respect to the ground truth, conditioned on the input image. By minimizing it, the generator tries to fool the discriminator.

– **Perceptual loss:**

$$l_{\text{perceptual}} = \frac{1}{L} \|F_{vgg-19}(I_{gt}) - F_{vgg-19}(I_g)\|_2, \quad (4)$$

where  $F_{vgg-19}(I) = \text{Concat}(REL U_2(I), REL U_3(I), REL U_4(I))$ ,  $REL U_i(I)$  is the feature after  $REL U$  activation function [47] in  $i$ th layer in VGG-19 network for image  $I$ , and  $L$  is the length of the concatenated feature. As with [48,49], we use the pre-trained VGG-19 network [50] to extract the high-level visual feature.

The final loss for the generator is a weighted combination of those three losses:

$$l_G = l_{\text{content}} + \alpha l_{\text{adversarial}} + \beta l_{\text{perceptual}}. \quad (5)$$

where  $\alpha = 1$  and  $\beta = 0.02$  based on tuning the validation set.

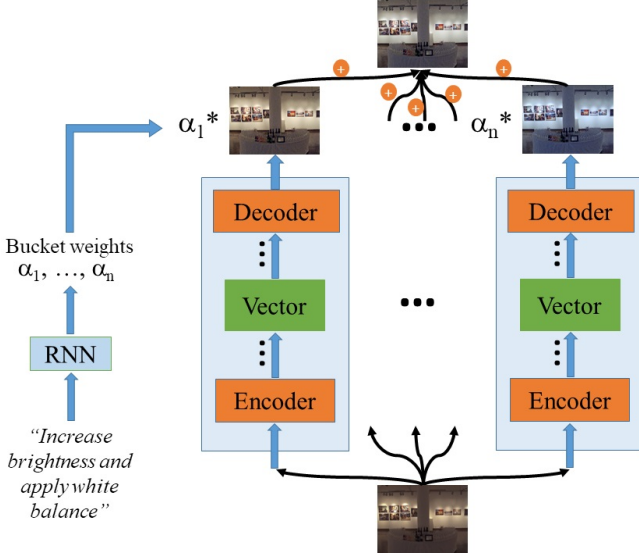


Fig. 5: Bucket model.

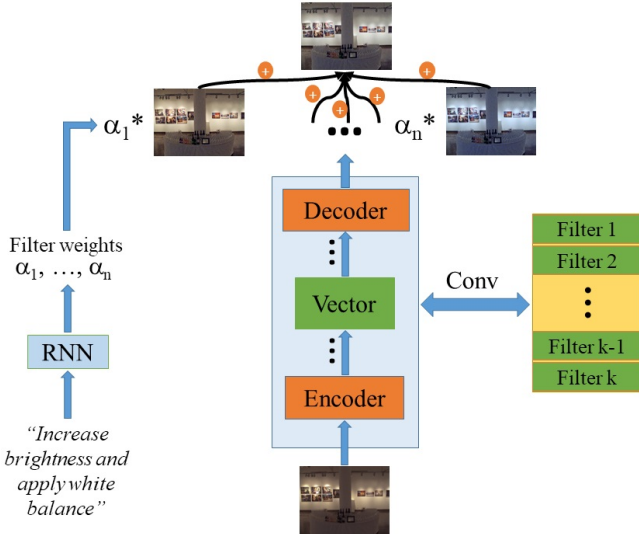


Fig. 6: Filter bank model.

**Bucket Model.** One design is the bucket model, which is based on the idea that similar image transformations should be grouped as buckets. Each bucket represents a different

image transformation (e.g., one for increasing the brightness, another for reducing the contrast). The disadvantage is the grouping is manual. The architecture of the bucket model is shown in Figure 5.

Given some text, we train an RNN to generate a distribution over buckets, and the final generated image is a weighted linear combination of different buckets. Let

$$h(T_{des}) = \overrightarrow{RNN}(t_1, \dots, t_n) || \overleftarrow{RNN}(t_1, \dots, t_n), \quad (6)$$

where  $\overrightarrow{RNN}$  and  $\overleftarrow{RNN}$  are the last hidden state vectors when text is fed to RNN in opposite directions. Let weight  $\alpha = \text{softmax}(h(T_{des}))$ , with  $K_b$  buckets and the output of each bucket being  $I_k$ . (In our work,  $K_b = 5$ .) The final output image  $I_g$  is a weighted linear combination from the different buckets, i.e.,  $I_g = \sum_{k=1}^{K_b} \alpha_k I_k$ .

In this model, each bucket has its own encoder-decoder architecture. The encoder is a down-sampling procedure which contains a series of conv-batch normalization [51,46] Leaky ReLU units [47]. The decoder is similarly constructed, except in reverse order to constitute an up-sampling procedure. In our implementation, the down-sampling and up-sampling networks have the same depth, and optionally we can use skip connection [52], i.e., we concatenate the  $i$ th layer in down-sampling network with the  $(N-i)$ th layer in up-sampling network.

To group the image, several methods can be used: surface form level word matching, cluster over word (sentence) embedding, or manually design the buckets. In our work, however, we manually designed the buckets based on the bigram distribution shown in Figure 22 in the supplementary file. We have tried using automatic grouping methods, but they appear to be less effective.

**End-to-End Model.** The bucket model, while straightforward, requires some hand-crafting of the buckets. Inspired by [33], we also design an end-to-end model. We use another RNN (which is different from that used in the discriminator) to encode the text to a vector; this vector is then concatenated with the image vector. As with the bucket model, we also use the encoder-decoder framework to encode the image. The overall architecture of the end-to-end model is shown in Figure 4.

Given an image  $I_{img}$ , we first encode it through a deep convolution neural network as  $\text{Encode}(I_{img})$ , followed by a depth concatenation between this image vector and the text vector:

$$h(T_{des}, I_{img}) = \text{DepthConcat}(\text{Encode}(I_{img}), h(T_{des})). \quad (7)$$

Subsequently, we feed  $h(T_{des}, I_{img})$  to the decoder.

**Filter Bank Model.** An end-to-end model is conceptually elegant. However, making it work is difficult due to limited expressive power, especially if we consider that image transformations can be bidirectional. An example is with respect to brightness, where the user can specify to “increase the brightness” or “decrease the brightness”. The bucket model is easy to understand, but it requires pre-designing the buckets; incorporating additional data will likely to require changes to the bucket design. Our third model, the filter bank model, is designed to combine the advantages of these two models. The architecture for the filter bank model is depicted in Figure 6.

Given a description, an RNN is used to encode it and generate a distribution over different filters, which is conceptually the same as the bucket model. Each filter  $F_k$  is a  $k \times k \times c_{in} \times c_{out}$  convolution filter, and the final image is a weighted linear combination of different images.

Given an image  $I_{img}$ , to generate the enhanced image based on filter  $F_k$ , we first use the encoder to encode the image as a hidden vector. We then convolve this hidden vector with filter  $F_k$ , and the result is fed to the decoder. With  $K_f$  filters, we have  $K_f$  different images generated. (In our work,  $K_f = 5$ .) For the  $k$ th filter, we have

$$I_k = \text{Decoder}(\text{Conv}(\text{Encode}(I_{img}), F_k)). \quad (8)$$

The final output image  $I_g$  is obtained using  $I_g = \sum_{k=1}^{K_f} \alpha_k I_k$ . This model is similar to that described in [53], but there is a major difference: The model in [53] is used for style transfer and each filter corresponds to one pre-determined style. During training, each training instance contains an image pair and corresponding filter id, and it only optimizes the corresponding filter and the shared encoder-decoder parameter. By comparison, for our filter bank model, the filters are jointly trained automatically from image pairs and the model learns how to decompose the transformation automatically (the only manual step is specifying the number of filters).

## 4 Data Collection

To train our models, we need original-edited image pairs with associated text descriptions. To the best of our knowledge, there is no such existing dataset. The MIT-Adobe 5k dataset [10] consists of original-edited image pairs generated by five professional photographers, but it does not contain text that describe the image transformation (such as brightness change and color balance). For each image pair, the list of operations used to generate the edited image is given; an operation consists of a software editing command and its associated parameters. The fine granularity of information is not useful for associating general casual description of the image transformation with the original-edited image pair. Other publicly available text-image datasets such as MS-COCO [54], ReferIt [55], and Flickr30k Entities [56] contain text that describe the image content, but such text are not related to image editing or style. In addition, these datasets do not contain edited versions of the original.

As such, we ran a user study to collect our own dataset through Amazon Mechanical Turk. We use a random subset of the MIT-Adobe 5k dataset; for a given original-edited image pair, we ask the subject to type in a phrase to describe the image transformation. We also flip the order of the image pair to sample the reverse transformation.

Each task (“hit” in AMT parlance) involves describing transformations for 8 pairs. For each image pair, the subject was asked to rate the image transformation and describe the image operations that are applied to the original to produce the edited version.



Procuring reliable data from such a user study is difficult because most Turkers are not highly familiar with concepts of photography, and as such, have only rudimentary vocabularies to describe visual changes. Initially, to assist with the task, we provided several example image pairs with plausible responses as guidelines. This unfortunately



resulted in subjects copying and pasting example responses regardless of relevance. Even if they do not copy and paste responses, many users are not familiar with imaging concepts and provided inappropriate text.

In response to these issues, we made the following changes: (1) disabled copy and paste, (2) added examples (with explanations) that would cause their work to be rejected, (3) added a qualification test to see if the subject understands color and contrast, and (4) used heuristics to manually filter out “bad” responses. The new data are significantly better than those obtained through the trial run. By disabling cut-and-paste, the responses are much more varied. By explaining why responses may be rejected and enforcing a qualification test, data noise is significantly reduced.

The interface with examples is shown in Figure 7. (See Section A in the supplementary file for additional examples, the qualification test, and task interface.)


⇒


Original Image
Edited Image

**Rating:**

- ☐ ☆ : quality of edited image is much worse than the original
- ☐ ☆☆ : quality of edited image is worse than the original
- ☐ ☆☆☆ : quality of edited image is similar to the original
- ☐ ☆☆☆☆ : quality of edited image is better than the original
- ☒ ☆☆☆☆☆ : quality of edited image is much better than the original

**Describe the editing operation(s):**

**Acceptable Response for this Instance:**

- \* increase color saturation
- \* improve color balance
- \* improve white balance
- \* deepen the colors
- \* make it look less washed out
- \* make the colors stand out more

**Unacceptable Response for this Instance (reason for rejection):**

- \* increase color saturation, improve color balance, improve white balance and deepen the colors (too many operations, at most two is acceptable)

Fig. 7: Guidelines and example responses provided in the user study.

| Model                   | <i>p</i> -value      |
|-------------------------|----------------------|
| End2end vs. GT          | $7.4 \times 10^{-6}$ |
| GT vs. Bucket(a)        | 0.007                |
| End2end vs. Bucket(f)   | 0.02                 |
| GT vs. Bucket(f)        | 0.02                 |
| End2end vs. Bucket(a)   | 0.06                 |
| FB vs. End2end          | 0.09                 |
| FB vs. GT               | 0.09                 |
| FB vs. Bucket(f)        | 0.53                 |
| FB vs. Bucket(a)        | 0.58                 |
| Bucket(a) vs. Bucket(f) | 0.67                 |

Fig. 8: Pairwise comparison between different models, the lower the value of *p*, the more different the two models are.

Once the data have been collected, we further manually check the response. We removed responses that are obviously inconsistent with the actual image operation, too generic (e.g., “beautify the image”), or are not descriptions (e.g., “the edited image need to be brighten” in response to the edited image being a darkened version of the original). Totally 370 responses are removed in this way, which count 15% of all the raw responses.

We end up with 1884 image pairs and annotations, with each image pair having on average 1.6 text annotations. 1378 image pairs are used for training, 252 for validation and 252 for test. These image pairs contain multiple image transformation directions, e.g., improving the color balance, increase/decrease the image brightness, increase/reduce the image saturation, deepen the colors and keep the image the same. For additional statistics on the collected data, see Figs. 21 and 22 in the supplementary file.

## 5 Implementation

We use Pytorch to implement our models. We tried two different versions of encoder-decoder: one is a typical encoder-decoder without skip connection, with the other with skip connection [52,57]. We find the version with skip connection has better performance and faster training. We use Adam [58] as the optimizer, with the initial learning rate 0.001, and will half the learning rate if we don’t observe the loss reduction on validation set.

Our RNN is one layer bidirectional Gated Recurrent Unit (GRU) [59] with a hidden size of 128. The vocabulary size is around 4k and word embedding size is 200. We initialize the word embedding with the pre-trained Glove word embedding [60].

For our bucket and filter bank models, due to memory constraints, we limit the numbers of buckets and filters to 5 each, i.e.,  $K_b = K_f = 5$ . For our bucket model, we have  $K_b$  encoder-decoder pairs without any shared parameters. As a result, the optimization process requires a large amount of memory; in addition, it is slow, especially during back propagation. To overcome this problem, we pre-train the  $K_b$  independent

encoder-decoder pairs separately and fix them when training the bucket model. Additionally, for the bucket model, after training, we compute two image outputs, one with the highest weight (“argmax”, i.e., Bucket(a)) and another being a weighted average (“fusion”, i.e., Bucket(f)).

On the other hand, the filter bank and end-to-end models are trained from scratch, since their memory requirements are not as severe and the training is faster. Specifically, bucket model, including the parallel pre-training for different buckets, totally takes 40 hours and needs 4 GPUs, while filter bank takes 25 hours and only need 1 GPU, and the end-to-end model only needs 20 hours and 1 GPU. More implementation details are given in Section B in the supplementary file.

## 6 Experimental Results

In this section, we first report results for automatic image enhancement (without text) as a sanity check. We then describe the results of a user study to evaluate the performance of our models in producing the edited image given an input image and text description. Finally, we show the effects of the trained filters from our filter bank model.

### 6.1 Automatic Image Enhancement

We first investigate the performance of c-GAN with encoder-decoder architecture in the context of automatic image enhancement, without any text used. We randomly selected 1200 image pairs (results from one expert) from MIT-Adobe fiveK and train the model; a representative result on the validation set is shown in Figure 9. We quantitatively evaluate our automatic image enhancement performance. Table 1 lists the L2 error (in  $L^*ab$  space) for our method. Even though the randomly selected dataset in [11] is not the same with ours<sup>2</sup>, but in general, we believe our results are representative. The results indicate that c-GAN with encoder-decoder as generator is suitable for our problem.



Fig. 9: Examples of automatic image enhancement.

<sup>2</sup> The random dataset from [11] is not publicly available, so we instead randomly select the same number of images.

Table 1: Comparisons of average L2 error on test sets, with standard error of 95%.

|       | Input          | Hwang et al. [11] | Ours           |
|-------|----------------|-------------------|----------------|
| Error | $17.1 \pm 0.9$ | $15.0 \pm 0.8$    | $12.1 \pm 0.9$ |

## 6.2 Image Transformation from Text Description

Since there is no existing benchmark, we design a user study for such an evaluation. We are specifically interested in how well the edited image fit the text description given an input image, for all the models and ground truth<sup>3</sup>. We want to extract metrics that are both absolute (through standalone rating) and relative (pairwise comparison). One representative result is given in Figure 10.

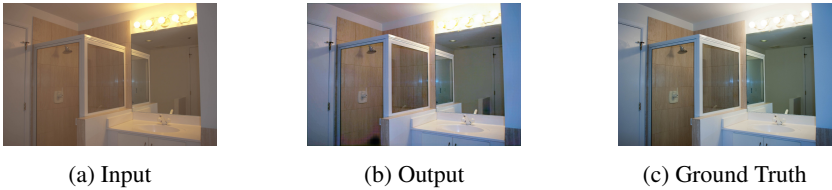


Fig. 10: Example of image editing under textual description “enhance white balance and contrast.” More examples are in the supplementary material.

**Standalone rating:** The subject is shown an original-edited image pair with text that describes the image transformation, and is asked to rate (on a scale of one to five stars) based on the instruction “how well does the edited image follow the instructions?”. There are five different pair versions, with the original image the same throughout and the edited image from ground truth, bucket model (fusion and argmax), filter bank model, and end-to-end model, respectively. The order of appearance is randomized. Each subject is shown eight image pairs corresponding to two different original images. For this portion of the user study, each image pair get five ratings, and the rating for that pair is averaged.

**Pairwise comparison:** The subject is shown two image pairs as well as the text description, and is asked to pick the pair that fits the text better. The same four versions are used. Each subject makes eight comparisons.

We obtained responses from 120 subjects; the reward for each task or “hit” is US\$0.20. (The user study interface is shown in Section A in the supplementary file.) Results of the user study are listed in Tables 2.

<sup>3</sup> Please note that we are less interested in how measuring how close the generated edited image is to the ground truth, because such a metric takes the text description out of the loop.

| Model        | Mean | Std Dev |
|--------------|------|---------|
| Ground Truth | 3.53 | 1.22    |
| Bucket(f)    | 3.36 | 1.25    |
| Bucket(a)    | 3.33 | 1.21    |
| Filter Bank  | 3.31 | 1.22    |
| End-to-End   | 3.19 | 1.30    |

Table 2: Standalone rating for the different models.

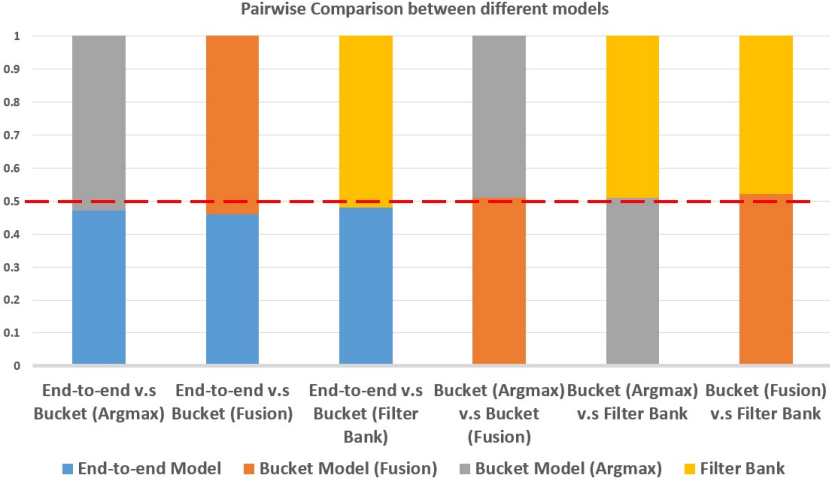


Fig. 11: Pairwise rating between different models. Red dash line represents equal rating 0.5.

Table 2 shows that the bucket model has the highest rating among all the three models. This is not surprising since the bucket model is customized, with the disadvantage of being less scalable. The filter bank model is next best with the end-to-end model being third. While the end-to-end model is the most conceptually elegant with the least amount of user specification, it has only one encoder-decoder, which limits its expressive power. It is less able to learn multiple directional transformations. While the filter bank model also has only one encoder-decoder, it has filters between them; the filters can be interpreted as a type of bucket model with shared encoder-decoder parameters among the buckets.

Compared with the ground truth, however, the differences are not significant. Please note that the ground truth version has a score of only 3.53; this may be due to most users being not familiar with image concepts. Figure 11 shows pairwise ratings between different models, which is consistent with Table 2.

Table 8 lists the  $p$ -values between the scores of different models, where smaller  $p$ -values implies larger difference between models<sup>4</sup>. Based on this table, the filter bank model is very close to the bucket model while the end-to-end model is less similar to the bucket model or filter bank model.

From the practical point of view, the filter bank model appears to be the best choice since the performance is good while not requiring much manual effort (apart from selecting the number of filters). Additionally, it requires less memory than the bucket model. For same encoder-decoder architecture with  $K_b$  buckets, it only need  $1/K_b$  memory as that for the bucket model. In addition, the filter bank model does not require pre-training for different buckets, making it much more efficient. For the same amount of memory, the filter bank model can afford to incorporate more filters than there are buckets.

### 6.3 Effects of Automatically Trained Filters

A significant advantage of the filter bank model is that we do not need to manually design the filters. In this section, we show some results of applying the automatically trained filters. Interestingly, each filter appears to correspond to a specific transformation. For example, the filter  $F_1$  corresponds to brightness reduction while filter  $F_2$  corresponds to brightness increase. This is consistent with [53], except that we do not explicitly specify each filter’s function. Figure 12 shows images generated by different filters.

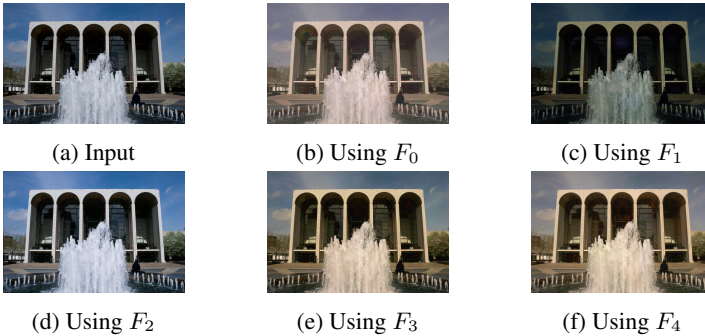


Fig. 12: Effect of different filters.

### 6.4 Observations on Learned Transformation

Even though our models were trained on global tonal adjustments, they are able to learn local transformations. The RGB remapping distributions in Figure 13 for two representative images show that our transformation, unlike its counterpart for expert A

<sup>4</sup> For a given model, we calculate the average score for each image pair and then evaluate the  $p$ -values between the scores of different models.

in the MIT-Adobe 5k dataset, is local. This is evident from the significantly more spread out distributions; for our method, each RGB input is mapped to wider range of outputs compared with the expert A. This results demonstrate that our models learn much more complicated mapping other than a single global mapping.

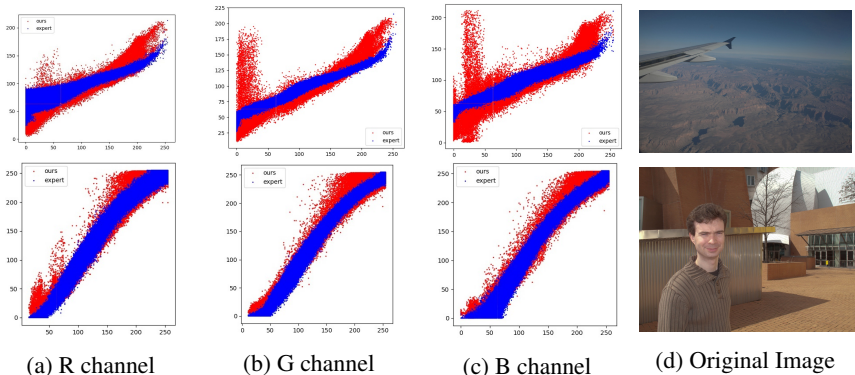


Fig. 13: RGB remapping distributions. For the expert enhanced image (expert A in MIT-Adobe 5k dataset), the mappings are almost one-to-one (in blue), while those for our edited images (in red) are not, demonstrating our editing is local.

## 6.5 Using Graph RNN on Filter Bank Model

The filter bank model is a good compromise between manual effort and performance. To further investigate the influence of text encoding, we also replaced RNN with Graph RNN, more specifically, Graph GRU (Gated Recurrent Unit). The graph structure is obtained from [44]; the last hidden state of graph RNN in both directions are used to represent the text. Please note that that conventional RNN is still used in the discriminator. Table 3 shows that Graph GRU performed better than plain GRU. One explanation is Graph GRU can utilize the dependency structures between different tokens and ignore the less important words in textual instructions. We also found that the Graph GRU is more effective when the text description is long and ambiguous. For more analysis with Graph RNN, see Section D in the supplementary file.

Table 3: Performance comparison between different RNNs in our filter bank model.

|                   | Graph GRU       | GRU             |
|-------------------|-----------------|-----------------|
| Standalone rating | $3.35 \pm 1.24$ | $3.31 \pm 1.22$ |
| Pairwise rating   | 0.52            | 0.48            |

## 7 Concluding Remarks

We show how we train a system to globally edit an image given a general textual command. To this end, we propose three models (bucket, filter bank, and end-to-end), which have different requirements in terms of initialization, memory requirements, and amount of training. Given the lack of database on image pair with text descriptions, we collected one on our own. Experimental results validate our models, and we believe our work is the first to address the general computational photography application of editing images purely through textual description.

One current limitation is we handle only editing based on global transformations (even the learned the transformation is local). To allow object-based editing, we would need to integrate object segmentation with a natural language module [7,36] to our system or design a joint model which can simultaneously segment and transform [36]. At the same time, we found it's difficult to obtain large scale, while diverse enough data, one possible way to alleviate this issue is data augmentation [61]. Given the facts user usually prefer a series of simple, consecutive and coherent textual description, another interesting direction is to extend our work in chat environment [29,62]. Finally, we can also investigate the multi-DSSM [63] loss in our model [35]. We leave the development of all these functionality for future work and we believe this will be an exciting and important research topic.



## References

1. Laput, G., Dontcheva, M., Wilensky, G., Chang, W., Agarwala, A., Linder, J., Adar, E.: Pixeltone: A multimodal interface for image editing. *Human-Computer Interaction International Conference (HCI)* (2013)
2. Cheng, M.M., Zheng, S., Lin, W.Y., Vineet, V., Sturges, P., Croo, N., Mitra, N., Torr, P.: Imagespirit: Verbal guided image parsing. *ACM Transactions on Graphics* (2014)
3. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. *Neural Information Processing Systems (NIPS)* (2014)
4. Kindermann, R., Snell, J.L.: Markov random fields and their applications. *American Mathematical Society* (1980)
5. Chen, J., Shen, Y., Gao, J., Liu, J., Liu, X.: Language-based image editing with recurrent attentive models. *arXiv preprint arXiv:1711.06288* (2017)
6. Seitaro, S., Koichiro, Y., Sakriani, S., Yu, S., Satoshi, N.: Interactive image manipulation with natural language instruction commands. *arXiv preprint arXiv:1802.08645* (2018)
7. Hu, R., Rohrbach, M., Darrell, T.: Segmentation from natural language expressions. *European Conference on Computer Vision (ECCV)* (2016)
8. Luong, T., Pham, H., Manning, C.D.: Effective approaches to attention-based neural machine translation. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. (2015) 1412–1421
9. Peng, N., Poon, H., Quirk, C., Toutanova, K., Yih, W.t.: Cross-sentence n-ary relation extraction with graph lstms. *Transactions of the Association of Computational Linguistics* **5**(1) (2017) 101–115
10. Bychkovsky, V., Paris, S., Chan, E., Durand, F.: Learning photographic global tonal adjustment with a database of input / output image pairs. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2011)
11. Hwang, S., Kapoor, A., Kang, S.B.: Context-based automatic local image enhancement. *European Conference on Computer Vision (ECCV)* (2012)
12. Kapoor, A., Caicedo, J.C., Lischinski, D., Kang, S.B.: Collaborative personalization of image enhancement. *International Journal of Computer Vision (IJCV)* (2013)
13. Yan, J., Lin, S., Kang, S.B., Tang, X.: A learning-to-rank approach for image color enhancement. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2014)
14. Yan, Z., Zhang, H., Wang, B., Paris, S., Yu, Y.: Automatic photo adjustment using deep neural networks. *ACM Transactions on Graphics* (2015)
15. Gatys, L., Ecker, A., Bethge, M.: Image style transfer using convolutional neural networks. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
16. Hwang, Y., Lee, J.Y., Kweon, I.S., Kim, S.J.: Color transfer using probabilistic moving least squares. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2014)
17. Lee, J.Y., Sunkavalli, K., Lin, Z., Shen, X., Kweon, I.S.: Automatic content-aware color and tone stylization. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
18. Liu, Y., Cohen, M., Uyttendaele, M., Rusinkiewicz, S.: Autostyle: Automatic style transfer from image collections to users' images. *Eurographics* (2014)
19. Vinyals, O., Toshev, A., Bengio, S., Erhan, D.: Show and tell: A neural image caption generator. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)
20. Chen, X., Zitnick, C.L.: Mind's eye: A recurrent visual representation for image caption generation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)
21. Donahue, J., Darrell, T.: Long-term recurrent convolutional networks for visual recognition and description. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)

22. Liu, C., Mao, J., Sha, F., Yuille, A.: Attention correctness in neural image captioning. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
23. Dai, B., Lin, D., Urtasun, R., Fidler, S.: Towards diverse and natural image descriptions via a conditional gan. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
24. Huang, T.H.K., Parikh, D., Vanderwende, L., Galley, M., Mitchell, M.: Visual storytelling. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
25. Venugopalan, S., Xu, H., Saenko, K.: Translating videos to natural language using deep recurrent neural networks. *International Conference on Computer Vision (ICCV)* (2015)
26. Antol, S., Agrawal, A., Batra, D., Zitnick, C.L., Parikh, D.: Vqa: Visual question answering. *International Conference on Computer Vision (ICCV)* (2015)
27. Yang, Y., Li, Y., Fermuller, C., Aloimonos, Y.: Neural self talk: Image understanding via continuous questioning and answering. *Arxiv* (2015)
28. Mostafazadeh, N., Misra, I., Devlin, J., Mitchell, M., He, X., Vanderwende, L.: Generating natural questions about an image. *Annual Meeting of the Association for Computational Linguistics (ACL)* (2016)
29. Das, A., Kottur, S., Gupta, K., Singh, A., Yadav, D., Moura, J.M.F., Parikh, D., Batra, D.: Visual dialog. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
30. Socher, R., Karpathy, A., Le, Q.V., Manning, C.D., Ng, A.Y.: Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics (TACL)* (2013)
31. Hu, R., Xu, H., Rohrbach, M., Feng, J., Saenko, K., Darrell, T.: Natural language object retrieval. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
32. Rohrbach, A.: Grounding of textual phrases in images by reconstruction. *European Conference on Computer Vision (ECCV)* (2016)
33. Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text-to-image synthesis. *International Conference on Machine Learning (ICML)* (2016)
34. Yan, X., Yang, J., Sohn, K., Lee, H.: Attribute2image: Conditional image generation from visual attributes. *European Conference on Computer Vision (ECCV)* (2016)
35. Xu, T., Zhang, P., Huang, Q., Zhang, H., Gan, Z., Huang, X., He, X.: Attngan: Fine-grained text to image generation with attentional generative adversarial networks. *arXiv preprint arXiv:1711.10485* (2017)
36. Hong, S., Yang, D., Choi, J., Lee, H.: Inferring semantic layout for hierarchical text-to-image synthesis. *arXiv preprint arXiv:1801.05091* (2018)
37. Mao, J.: Generation and comprehension of unambiguous object descriptions. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
38. Yu, L., Tan, H., Bansal, M., Berg, T.L.: A joint speaker-listener-reinforcer model for referring expressions. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
39. Luo, R., Shakhnarovich, G.: Comprehension-guided referring expressions. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
40. Karpathy, A., Joulin, A., Li, F.F.: Deep fragment embeddings for bidirectional image sentence mapping. *Conference on Neural Information Processing Systems (NIPS)* (2014)
41. Karpathy, A., Li, F.F.: Deep visual-semantic alignments for generating image descriptions. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)
42. Kong, C., Lin, D., Bansal, M., Urtasun, R., Fidler, S.: What are you talking about? text-to-image coreference. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2014)
43. Tai, K.S., Socher, R., Manning, C.D.: Improved semantic representations from tree-structured long short-term memory networks. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Volume 1. (2015) 1556–1566

44. Manning, C., Surdeanu, M., Bauer, J., Finkel, J., Bethard, S., McClosky, D.: The stanford corenlp natural language processing toolkit. In: Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations. (2014) 55–60
45. Mirza, M., Osindero, S.: Conditional generative adversarial nets. Arxiv (2014)
46. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Instance normalization: The missing ingredient for fast stylization. Arxiv (2016)
47. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. IEEE Conference on Machine Learning (ICML) (2010)
48. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. European Conference on Computer Vision (ECCV) (2016)
49. Ledig, C., Theis, L., Huszar, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., Shi, W.: Photo-realistic single image super-resolution using a generative adversarial network. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
50. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. International Conference on Learning Representations (ICLR) (2014)
51. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. International Conference on Machine Learning (ICML) (2015)
52. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
53. Chen, D., Yuan, L., Liao, J., Yu, N., Hua, G.: Stylebank: An explicit representation for neural image style transfer. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
54. Lin, T.Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P.: Microsoft coco: Common objects in context. European Conference on Computer Vision (ECCV) (2014)
55. Kazemzadeh, S., Ordonez, V., Matten, M., Berg, T.L.: Referit game: Referring to objects in photographs of natural scenes. Conference on Empirical Methods in Natural Language Processing (EMNLP) (2014)
56. Plummer, B.A.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. International Journal of Computer Vision (IJCV) (2016)
57. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
58. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. International Conference on Learning Representations (ICLR) (2014)
59. Chung, J., Gulcehre, C., Cho, K.H., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. NIPS 2014 Workshop on Deep Learning (2014)
60. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. Empirical Methods in Natural Language Processing (EMNLP) (2014)
61. Dong, H., Zhang, J., McIlwraith, D., Guo, Y.: I2t2i: Learning text to image synthesis with textual data augmentation. arXiv preprint arXiv:1703.06676 (2017)
62. Sharma, S., Suhubdy, D., Michalski, V., Kahou, S.E., Bengio, Y.: Chatpainter: Improving text to image generation using dialogue. arXiv preprint arXiv:1802.08216 (2018)
63. Huang, P.S., He, X., Gao, J., Deng, L., Acero, A., Heck, L.: Learning deep structured semantic models for web search using clickthrough data. In: Proceedings of the 22nd ACM international conference on Conference on information & knowledge management, ACM (2013) 2333–2338

In this supplementary file, we provide more information on the interface used for our data collection (Section 1) and on our implementation (Section 2). We also show additional results for our three models (Section 3).

## A Interface for Data Collection

The interface used for our data collection (through Amazon Mechanical Turk) consists of the following parts:

- Introduction page that lists instructions (Figure 14),
- Examples as guidance (Figure 15),
- Qualification test to ensure the subject has some understanding of image concepts (Figure 16), and
- Data collection (Figures 17).

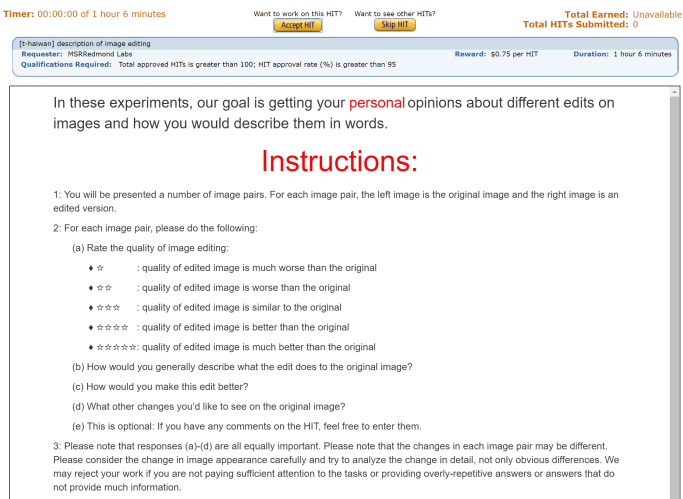


Fig. 14: Instructions for data collection.

The interface used for evaluation consists of:

- Instructions (Figure 18),
- Examples as guidance (Figure 19), and
- Rating (Figure 20).

We analyzed the data collected through unigram statistics (Figure 21) and bigram statistics (Figure 22). Please note that the orange curve represents cumulative frequency.

How would you improve the editing:

**Acceptable Response for this Instance:**

- \* make grass greener
- \* make skin color less red
- \* make the sky more blue

**Unacceptable Response for this Instance (reason for rejection):**

- \* remove the car (editing shouldn't change the image content)
- \* change the sky color (too vague)
- \* add haze to the image (this will worsen the image)

What other editing to the original image would you suggest:

**Acceptable Response for this Instance:**

- \* blur background slightly
- \* increase contrast of people
- \* increase brightness of image

**Unacceptable Response for this Instance (reason for rejection):**

- \* reduce the brightness (this will worsen the image)
- \* make the people less happy (only color/intensity change is allowed)
- \* reduce the sharpness (this will worsen the image)

Fig. 15: Examples provided as guidelines.

Please pass the following qualification test first:

If we want to make the image darker, we need to adjust its brightness

☐ correct ☐ wrong

If we want to make a certain region stand out, we might adjust its contrast

☐ correct ☐ wrong

The sharpness of blurry image generally is high

☐ correct ☐ wrong

Deepening the colors is the same as increasing brightness

☐ correct ☐ wrong

The foreground or background of an image can contain several objects

☐ correct ☐ wrong

White balance is different from color balance

☐ correct ☐ wrong

Start

Fig. 16: Qualification test.

## B Implementation Details

The convolution kernels used are  $4 \times 4$  spatial filters with stride 2. The downsample ratio is 2 for both encoder and discriminator, while the upsample ratio is 2 for the decoder. We find the skip connection helps to speed up the training and improve the editing performance. With the skip connections used, we have 8 Convolution-InstanceNorm-Leaky-ReLU layers. The number of filters in each layer in encoder are 64-128-256-512-512-512-512, while in decoder, they are 512-1024-1024-1024-1024-512-256-128. The filter numbers in discriminator are 64-128-256-512. We use the default Leaky-ReLU function without any parameter modification.

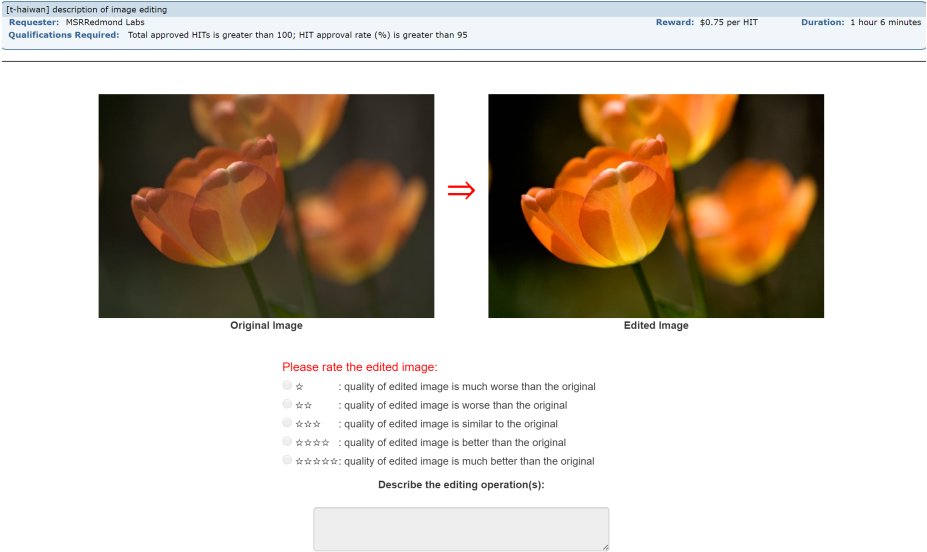


Fig. 17: Interface for rating the quality of the edited image relative to the original and for providing text description.

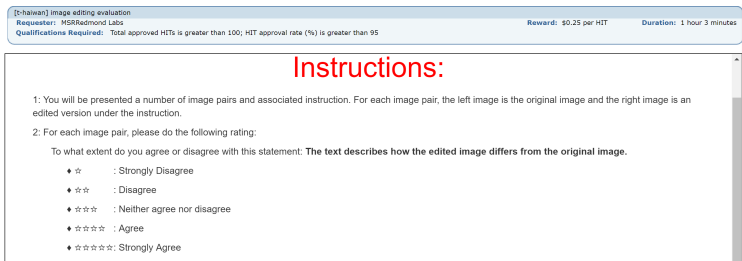


Fig. 18: Introduction for rating the image editing under the text description.

For the RNN, we use a one-layer bi-directional GRU. All the models are trained 100 epochs, with each epoch taking around 20 minutes on a single TitanX GPU.

We have tried a number of designs for the buckets:

- Surface form bag-of-word cluster over the descriptions,
- Cluster over the descriptions use sentence embedding obtained from various ways:
  - Averaged on all tokens’ embedding,
  - Element-wise product on all tokens’ embedding, and
  - Element-wise max on all tokens’ embedding,
- Manual design.

We found the manually designed the bucket to produce the best results. To design the buckets, we looked at the bigram statistics (Figure 22) and images in the dataset to select major image attributes, such as brightness, contrast, and white balance. We then

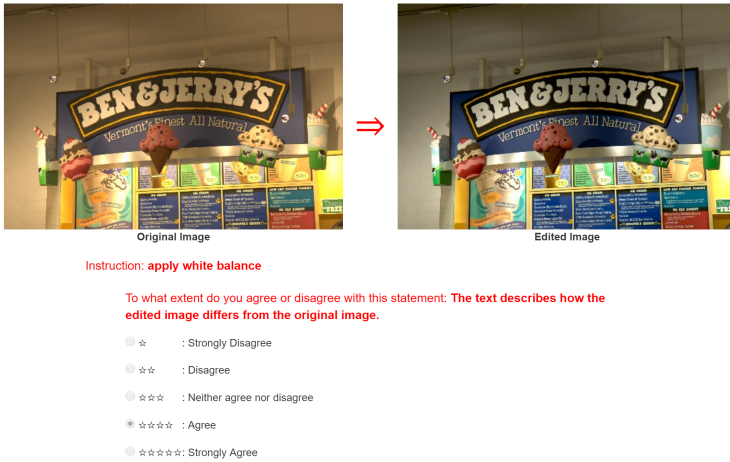


Fig. 19: Examples provided for rating the image editing under the text description.

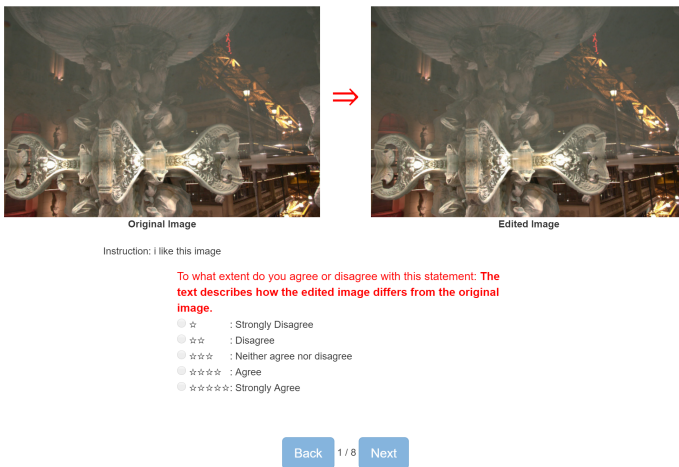


Fig. 20: Interface for rating the image editing under the text description

partition the images into groups, each to be trained as a bucket, using image statistics such as difference in mean gray value and mean RGB vector between the image pairs. Please note that an image may belong in more than one group/bucket.

Initially, we found that the generator is capable of handling changes in only one direction; for example, if we train a bucket to make the image brighter, then no matter what kind of image is given, the model will always try to make the image brighter, even if the image brightness is already high. To alleviate this problem, we augment the training data: we use the same image as input and ground truth to make a new pair and we manually generate some description according to some templates as following:

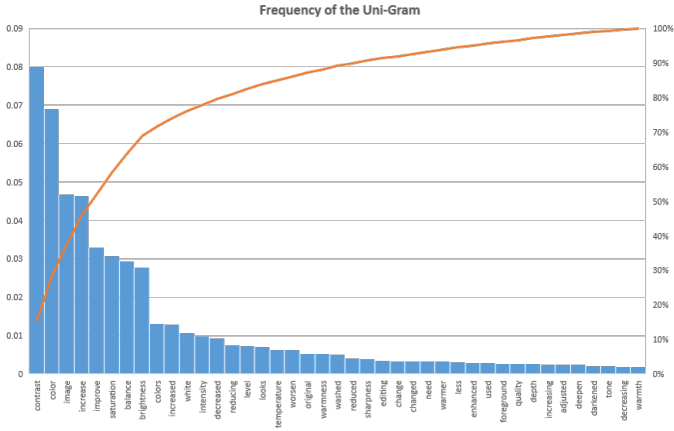


Fig. 21: Unigram statistics of descriptions.

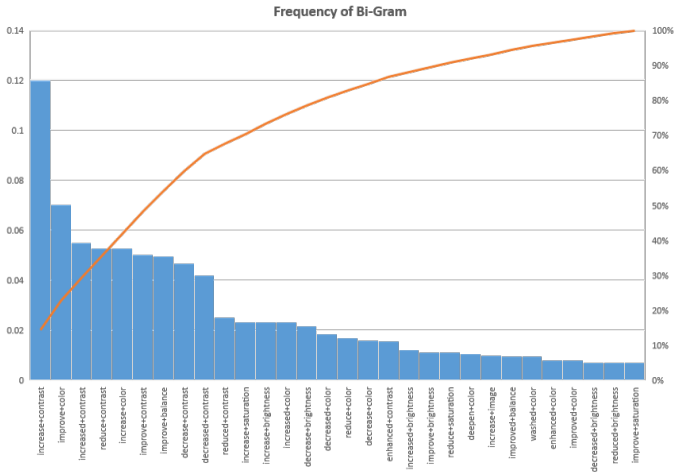


Fig. 22: Bigram statistics of descriptions.

- it is good; the image is good; this picture is amazing; this picture looks good;
- I would like to share this image; I would like to send this image to my friend;
- the tone in this image is good; the tone in this image is perfect;

we can replace the underlined text with any paraphrases.

To sample a random text for the discriminator (Section 3.1 in the paper), we first calculate the confusion matrix between most frequent unigram (Figure 21) by checking their co-occurrence in the description. Given a description, we first find the most unrelated unigrams and then find a description that contains these unigrams.



## C Additional Results

In this section, we show more results for generation of edited images given an input image and text description using our three models: bucket model, fusion version (Figure 23), bucket model, argmax version (Figure 24), filter bank model (Figure 25), and end-to-end model (Figure 26).

## D Graph GRU Example

More examples with Graph RNN is given in Figure 27.



(a) Input + “decrease saturation”



(b) Output



(c) Ground Truth



(d) Input + “improve color balance”



(e) Output



(f) Ground Truth



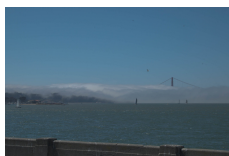
(g) Input + “increase saturation”



(h) Output



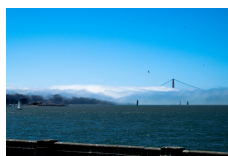
(i) Ground Truth



(j) Input + “color balance improved, sharpness improved”



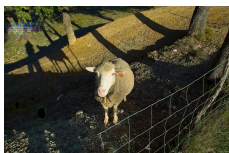
(k) Output



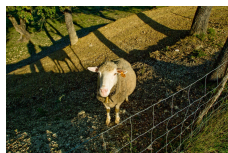
(l) Ground Truth



(m) Input + “saturation raised; colors are much deeper”



(n) Output



(o) Ground Truth



(p) Input + “color increased and skin tone of human is also good”



(q) Output



(r) Ground Truth

Fig. 23: More results for our bucket (fusion) model.



(a) Input + “decrease saturation”



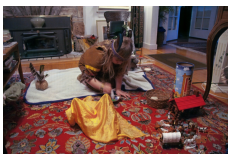
(b) Output



(c) Ground Truth



(d) Input + “improve color balance”



(e) Output



(f) Ground Truth



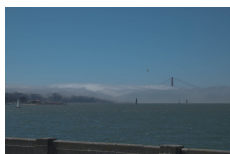
(g) Input + “increase saturation”



(h) Output



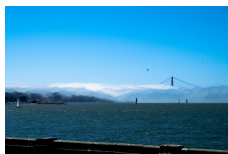
(i) Ground Truth



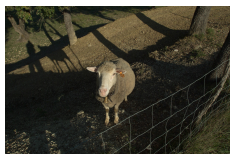
(j) Input + “color balance improved, sharpness improved”



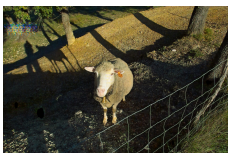
(k) Output



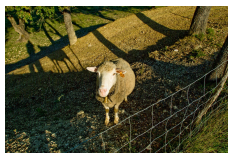
(l) Ground Truth



(m) Input + “saturation raised; colors are much deeper”



(n) Output



(o) Ground Truth



(p) Input + “color increased and skin tone of human is also good”



(q) Output



(r) Ground Truth

Fig. 24: More results for our bucket (argmax) model.



(a) Input + “decrease saturation”



(b) Output



(c) Ground Truth



(d) Input + “improve color balance”



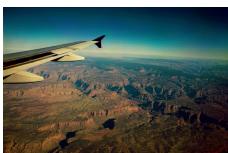
(e) Output



(f) Ground Truth



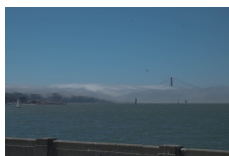
(g) Input + “increase saturation”



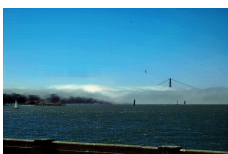
(h) Output



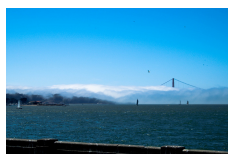
(i) Ground Truth



(j) Input + “color balance improved, sharpness improved”



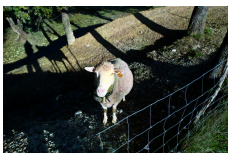
(k) Output



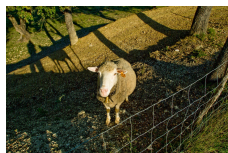
(l) Ground Truth



(m) Input + “saturation raised; colors are much deeper”



(n) Output



(o) Ground Truth



(p) Input + “color increased and skin tone of human is also good”



(q) Output



(r) Ground Truth

Fig. 25: More results for our filter bank model.





(a) Input + “decrease saturation”



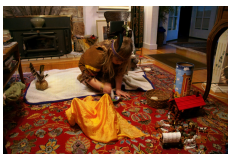
(b) Output



(c) Ground Truth



(d) Input + “improve color balance”



(e) Output



(f) Ground Truth



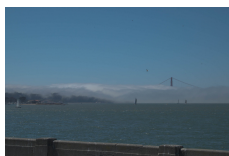
(g) Input + “increase saturation”



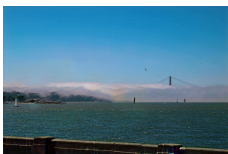
(h) Output



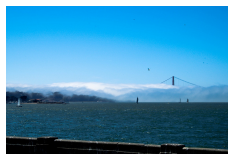
(i) Ground Truth



(j) Input + “color balance improved, sharpness improved”



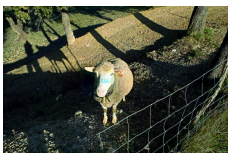
(k) Output



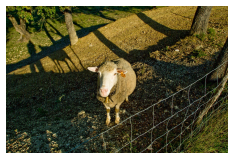
(l) Ground Truth



(m) Input + “saturation raised; colors are much deeper”



(n) Output



(o) Ground Truth



(p) Input + “color increased and skin tone of human is also good”



(q) Output



(r) Ground Truth

Fig. 26: More results for our end-to-end model.



(a) Input + “it looked as if they lowered the sharpness of the picture , as well as the brightness and the contrast .”



(b) Output with RNN



(c) Output with Graph RNN



(d) Input + “contrast is lowered . fading is used in this edited image . they decrease the sharpness , warmth and shadows .”



(e) Output with RNN



(f) Output with Graph RNN



(g) Input + “the editor made the mountains pop out more .”



(h) Output with RNN



(i) Output with Graph RNN

Fig. 27: More results for our filter bank model with Graph RNN.