LUCSS: Language-based User-customized Colorization of Scene Sketches

CHANGQING ZOU*, University of Maryland, College Park, United States

HAORAN MO*, Sun Yat-sen University, China

RUOFEI DU, University of Maryland, College Park, United States

XING WU, Sun Yat-sen University, China

CHENGYING GAO, Sun Yat-sen University, China

HONGBO FU, City University of Hong Kong, Hongkong

It's a moonlit night. There is a moon in the sky. A house is in the middle. A car is in front of the house. Three trees are on the left of the house. Two trees are on the right back of the car. There is a road.

It is a moonlit night. There is a **yellow** moon in the sky. An **orange** house with **light gray** roof is in the middle. A **yellow** car with **blue** window is in front of the house. Three **green** trees are on the left of the house. Two **green** trees are on the right back of the car. There is a **yellow** road. The sky is **blue** and all things are on **green** grass.

It is a moonlit night. There is a **purple** moon in the sky. A **yellow** house with **yellow** roof is in the middle. A **blue** car with **light gray** window is in front of the house. Three **yellow** trees are on the left of the house. Two **yellow** trees are on the right back of the car. There is a **light brown** road. The sky is **gray** and all things are on **yellow** grass.

It is a moonlit night. There is a **yellow** moon in the sky. A **blue** house with **blue** roof is in the middle. A **red** car with **blue** window is in front of the house. Three **green** trees are on the left of the house. Two **green** trees are on the right back of the car. There is a **dark gray** road. The sky is **blue** and all things are on **green** grass.

Fig. 1. We present LUCSS, a language-based interactive colorization system for scene sketches. It takes advantage of both instance level segmentation and language models in a unified generative adversarial network, allowing users to accomplish different colorization goals in a form of language instructions. Left: input scene sketch and its content description automatically generated by LUCSS. Three right columns show the colorization results generated by LUCSS following three different instructions at the bottom. Texts underlined are user-specified, with the target colors highlighted in bold.

Abstract. We introduce LUCSS, a language-based system for interactive colorization of scene sketches, based on their semantic understanding. LUCSS is built upon deep neural networks trained via a large-scale repository of scene sketches and cartoon-style color images with text descriptions. It consists of three sequential modules. First, given a scene sketch, the segmentation module automatically partitions an input sketch into individual object instances. Next, the captioning module generates the text description with spatial relationships based on the instance-level segmentation results. Finally, the interactive colorization module allows users to edit the caption and produce colored images based on the altered caption. Our experiments show the effectiveness of our approach and the desirability of its components to alternative choices.

CCS Concepts: • **Computing methodologies** → **Image Processing**;

Additional Key Words and Phrases: Deep Neural Networks, Sketch Captioning, Sketch Colorization, Scene Sketch

---

*Both authors contributed equally to the paper

1 INTRODUCTION

Sketching is one of the most efficient and compelling ways to communicate complex ideas among humans. While abstract sketches can be easily understood by us, teaching machines to understand the underlying semantics of sketches remains a challenging task. Recent research has achieved semantic understanding for individually-sketched objects and fostered applications such as sketch-based shape retrieval [Wang et al. 2015] and sketch classification [Eitz et al. 2012a].

Nevertheless, instance-level understanding, as often used in scene sketches, has not received much attention. Scene sketches usually contain multiple sketched objects, depict scenes of real or imaginary worlds, and widely appear in various scenarios, such as story books, sketch movies, and Computer-Aided Design (CAD) software. In this paper, we investigate the instance-level segmentation and interpretation of scene sketches. Our work can benefit a number of applications such as sketch-based co-retrieval [Xu et al. 2013] and context-based sketch classification [Zhang et al. 2018], which

currently take as input manual segmentation of individual scene objects.

With recent advances of deep neural networks and the availability of large-scale datasets such as MS COCO [Lin et al. 2014] and ImageNet [Deng et al. 2009], machines have outperformed humans in various image understanding tasks for natural images, such as image classification, face recognition, and image generation. Nevertheless, the power of deep networks for scene sketches is still unexplored, and the applications based on scene sketch understanding have rarely been investigated.

To address these problems, in [Zou et al. 2018], we have built a large-scale scene sketch dataset, called *SketchyScene*, and conducted initial studies on category-level semantic segmentation of scene sketches. In this paper, we investigate two research goals: 1) the capability of deep neural networks for scene sketch understanding, and 2) the potential applications based upon *instance-level* understanding.

To achieve these two goals in a unified framework, we present LUCSS, a language-based interactive colorization system, which consists of three interrelated modules: instance segmentation, captioning, and colorization. The instance segmentation module addresses how to employ deep networks for segmenting an input scene sketch into object instances. The colorization module can be considered as an application which produces a colorized image conforming to user-specified color requirements for segmented objects. Serving as a link between the segmentation and colorization modules, the captioning module takes the output of the segmentation module as input, and automatically generates a caption describing the input scene sketch. Motivated by the recent success of intelligent personal assistants such as *Apple Siri* and *Amazon Alexa*, which are enabled by speech recognition and natural language processing, we present a language-based approach that enables users to embed customized colors into the text description, instead of drawing color strokes on the sketch to specify the colors. Our approach is more compatible with voice commands for future multimodal colorization systems that can further enhance the user experience.

In this paper, we mainly address two challenging research problems. First, how shall we achieve precise instance-level segmentation? High-quality segmentation results are crucial to the subsequent colorization process, since users might specify different colors for individual objects. Unlike natural images, an input scene sketch consists of merely black lines and white background. Inferring instance segmentation of a sketched scene is thus challenging due to the sparsity of the visual features (for example, the foreground pixels only occupy 13% of all the pixels in *SketchyScene* [Zou et al. 2018]). Employing the state-of-the-art segmentation methods designed for natural images (e.g., [He et al. 2017]) directly on sketches does not provide promising results, as shown in Section 6. To address this problem, we enhance powerful segmentation models designed for natural image segmentation with the unique characteristics of scene sketches.

Second, how shall we colorize a high-resolution scene sketch with respect to language-based color inputs? Although the colorization of a single sketched object has been extensively investigated, the colorization of scene sketches with customized color labels remains an open problem. This challenge requires our system to build

accurate correspondences between the object instances (parts of object instances) and text-based color specifications. Additionally, it requires the system to infer both object-level and object-part-level segmentation. For example, a user may assign different colorization goals to the window of the car as in shown in Figure 1. To tackle this problem, we present a novel architecture which embeds LSTM (Long Short-Term Memory) to a sophisticated Generative Adversarial Network (GAN).

Furthermore, generating high-quality and high-resolution colorization results (768×768) is not a trivial task. We address this issue by using a two-stage pipeline consisting of object colorization and background colorization. Our experimental results (Section 7 and the supplementary materials) show that the LUCSS colorization framework achieves visually pleasing results.

We highlight the major contributions of LUCSS as follows:

(1) The first language-based, user-customizable colorization framework for scene sketches.
(2) The first solution for instance-level segmentation of scene sketches.
(3) A colorization dataset of scene sketches with text description and instance-level segmentation.

2 RELATED WORK

Our work is inspired by and build upon previous work in image segmentation, colorization, captioning, and generation with convolutional neural networks (CNNs) and conditional generative adversarial networks (cGANs).

2.1 Image Segmentation

In recent years, CNNs have been proven to yield the state-of-the-art accuracy in semantic object segmentation [Chen et al. 2017a, 2016; Shelhamer et al. 2016; Zhao et al. 2016]. The success of these methods is driven by large-scale, manually-annotated datasets, such as ImageNet [Deng et al. 2009] and MS COCO [Lin et al. 2014], which consist of millions of photographs with segmented objects. These methods often jointly predict a segmentation mask and an objectness score based on some appearance features that are specific to an individual class. Instance segmentation has become more common after the introduction of the R-CNN pipeline using category-independent region proposals [Hariharan et al. 2014]. Recently, a few methods have shown promise in the task of end-to-end trained instance segmentation [Girshick 2015; He et al. 2017; Ren et al. 2015a]. These approaches perform local and spatially-varying "objectness" estimates, with a simple global aggregation step in the end.

In our prior work [Zou et al. 2018], we conducted a pilot study on category-level semantic segmentation of scene sketches. In this paper, we investigate the problem of instance-level segmentation in scene sketches. Existing solutions for instance segmentation of natural images (e.g., [He et al. 2017]) cannot produce satisfactory results for scene sketches because they do not consider the unique characteristics of scene sketches. Please refer to Section 4 for further discussion.

2.2 Sketch Understanding

Sketch recognition is perhaps the most popular problem in sketch understanding. Since the debut of TU-Berlin dataset [Eitz et al.

2012b], many approaches have been proposed and the state-of-the-art approaches have even outperformed human beings in terms of the recognition accuracy [Yu et al. 2017]. Prior algorithms can be roughly classified into two categories: 1) those using hand-crafted features [Eitz et al. 2012b; Schneider and Tuytelaars 2014], and 2) those learning deep feature representations [Ha and Eck 2017; Yu et al. 2017]. The latter generally outperform the former by a clear margin.

Another stream of work has delved into parsing sketched objects into their semantic parts. Sun *et al.* [2012a] proposed an entropy descent stroke merging algorithm for both part-level and object-level sketch segmentation. Huang *et al.* [2014] leveraged a repository of 3D template models composed of semantically segmented and labeled components to derive part-level structures. Schneider and Tuytelaars [2016] performed sketch segmentation by looking at salient geometrical features (such as T-junctions and X-junctions) via a Conditional Random Field (CRF) framework. Instead of studying single object recognition or part-level sketch segmentation, we conduct an exploratory study for scene-level parsing of sketches, by using the large-scale scene sketch dataset SketchyScene [Zou et al. 2018].

2.3 Image Captioning

With the recent advances in deep neural networks and language learning models, a number of impressive algorithms for image captioning have been developed [Chen and Zitnick 2014; Donahue et al. 2014; Mao et al. 2015; Vinyals et al. 2016; Xu et al. 2015; Zhou et al. 2016]. We direct readers to a recent survey [Bernardi et al. 2016], which summarizes related datasets and an evaluation of various models. More recently, attention mechanisms have been applied to the network for narrowing down subjects and thus improving captioning results [Anderson et al. 2018; Das et al. 2017; Xu et al. 2015]. These modern image captioning models generally consist of two parts: an image encoder and a language model. The image encoder encodes a raw image into a feature map using a CNN, while the language model generates text sequentially with the extracted feature map and the inherited probabilistic dependency. In contrast to prior works which aim to generate human-like extractive captions, our captioning module directly produces a lower-level detailed description, which covers the entire scene sketch based on the results of instance segmentation. We use this instance-segmentation-based captioning module for user input assistance.

2.4 Scene Sketch Based Applications

While there is little work on semantic segmentation of scene sketches, some interesting applications have been proposed to utilize scenes with pre-tagged or pre-segmented sketched objects as input. For example, Sketch2Photo [Chen et al. 2009] combines sketching and photo montage for realistic image synthesis. Sketch2Cartoon [Wang et al. 2011] is a similar system which focuses on cartoon images. Similarly, Xu *et al.* [2013] propose Sketch2Scene, an automatic system which generates 3D scenes by co-retrieving and co-placing 3D shapes with respect to a scene of pre-segmented sketched objects. Sketch2Tag [Sun et al. 2012b] is a sketch-based image retrieval (SBIR) system, where scene items are automatically recognized and used as a text query to improve the retrieval performance. Our work provides automatic instance segmentation algorithms for scene sketches, and can immediately benefit the above applications.

2.5 Image Colorization

Image colorization assigns a three-dimensional label (RGB) to each pixel from an input gray-scale or sketch image. Early studies of colorization methods are mainly based on user interaction [Huang et al. 2005; Qu et al. 2006a; Yatziv and Sapiro 2006] or similar examples [Charpiat et al. 2008; Welsh et al. 2002] from gray-scale photographs. To achieve user-customized colorization, interactive strokes [Huang et al. 2005; Levin et al. 2004] are widely used to provide local guidance based on local intensity differences and spatial offsets [Luan et al. 2007; Qu et al. 2006b]. Further approaches using local guidance devise better similarity metrics by employing long range connections [An and Pellacini 2010; Xu et al. 2009] and local linear embeddings [Chen et al. 2005] to minimize user efforts. In addition to local guidance, non-local mean-based patch weight [Yao et al. 2010] and color palette [Chang et al. 2015] have also been proposed to provide global guidance.

Recent colorization systems [Cheng et al. 2015; Deshpande et al. 2015; Iizuka et al. 2016; Larsson et al. 2016; Yan et al. 2016], take advantage of deep CNNs and large-scale datasets to automatically produce plausible color images from gray-scale inputs. Adapting these models to scene sketches is a challenge since scene sketches are sparser than gray-scale images. Another line of research focuses on sketch colorization which generates color images from black-and-white sketches [Güçlütürk et al. 2016; Liu et al. 2017b; Sangkloy et al. 2017a; Xian et al. 2017]. However, most of the prior sketch colorization approaches target object-level sketches, while our work is a scene-level sketch colorization method.

Sangkloy *et al.* [2017b] has developed a system to translate sketches to real images, with colorization goals assigned by user color strokes. This system can well convey the user's colorization goals to image regions but its extensibility may be limited. PaintsChainer [Yonetsuji 2017] and Frans *et al.* [2017] have developed open-source interactive online applications for line-drawing colorization. Concurrently, Chen *et al.* [2017b] proposed a language-based colorization method for object-level sketches or gray-scale images. Our system is close to [Chen et al. 2017b] but we focus on scene-level sketch colorization.

2.6 Image Generation with GANs

Several recent studies, such as DCGAN [Radford et al. 2015], Pix2Pix [Isola et al. 2017], WGAN [Arjovsky et al. 2017], and SketchyGAN [Chen and Hays 2018], have demonstrated the value of variant GANs for image generation. The variations of conditional GANs have been further applied to text-to-image synthesis [Hong et al. 2018; Reed et al. 2016; Zhang et al. 2017], image inpainting [Pathak et al. 2016; van den Oord et al. 2016; Yeh et al. 2016], and image super-resolution [Ledig et al. 2017; Sønderby et al. 2016]. Nonetheless, all the aforementioned generators are conditioned solely on text or images. In contrast, LUCSS takes both image and text as input, presenting an additional challenge of fusing the features of a scene-level image and the corresponding text description.
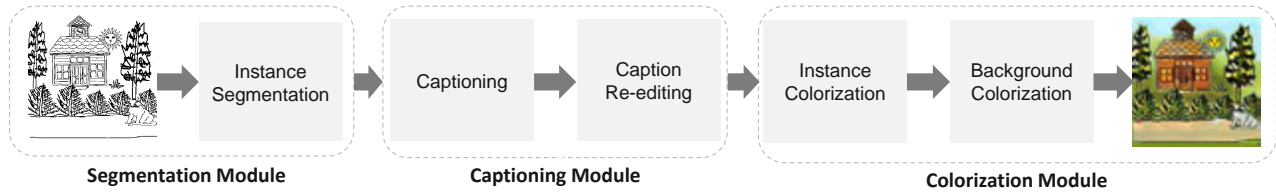
Fig. 2. The modules and the workflow of the LUCSS pipeline.

3 SYSTEM OVERVIEW

As illustrated in Figure 2, our system processes an input sketch through multiple stages. It first performs instance segmentation to recognize and locate individual objects in the scene sketch, and then automatically produces the caption describing lower-level detailed information of the sketch (*e.g.*, object category, object position, quantity) with a template-based algorithm. Next, the user can specify colorization goals by changing the automatically-generated caption. Finally, the system colorizes the whole sketch into a color image, where each segmented instance meets the user's requirements via a novel cGAN based model.

The system contains three sequential modules. The first one is an instance segmentation deep network. It takes as input a scene sketch image, and outputs the class labels and the instance identity label for each pixel belonging to the objects in the scene. In this module, we adapt the state-of-the-art segmentation models (including Deeplab V2 and Mask-RCNN) to the sketch data via analyzing the characteristics of the sketch data. The segmentation accuracy is significantly improved compared to the original models.

The second module is a language template-based caption generator. Taking the segmentation results as input, it analyzes the geometric relationships between object instances (e.g., position and occlusion relationships) and generates object-level scene descriptions. Figure 1 shows a typical result of this module. Besides captioning, this module provides a friendly interface for the user to assign different editing goals for individual object instances. This module can facilitate the system to build the exact correspondence between object instances in the scene and their corresponding sentences (since the caption is generated by the system from the segment instances, correspondences can be easily obtained by comparing the user-modified caption with the original one). In contrast to recent research [Chen et al. 2017b], which implicitly infers the correspondence between object parts and sentences using an attention mechanism, we leverage the captioning module to achieve more accurate correspondence, especially for complex scenes. As shown in the experiment of Section 7.1, implicit inference usually fails on a long caption with more than six sentences, but our approach has no such constraint. Hence it provides a better assurance that results will achieve user specifications.

The last module is a typical application based on the instance segmentation results. In this work we focus on the colorization task. Unlike recent research which aims to colorize single objects or relatively simple scenes [Chen et al. 2017b; Liu et al. 2017c; Varga et al. 2017], our work aims to solve a complex scene sketch colorization problem with the help of language-based user instruction. Our experiment in Section 5 found that even the most best model to date [Chen et al. 2017b] is unable to successfully colorize each
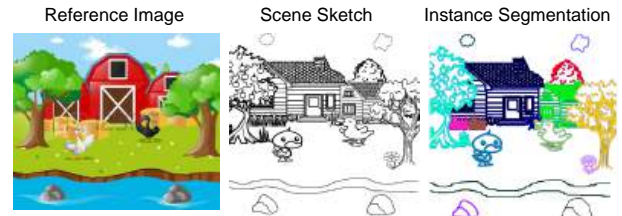

Fig. 3. An example scene template from SketchyScene. (a) shows a cartoon-style reference image harvested from the Internet, (b) shows a user-synthesized scene sketch using a repository of instance sketch, and (c) shows the ground-truth labels of the instance segmentation.

individual object instance to exactly meet the user's needs. The main reason is that a user's language instructions usually have ambiguity, especially when the target scene is complex (e.g., there are multiple object instances belonging to the same category in a scene). Our solution is to decompose the colorization of the whole scene sketch into two sequential steps: instance colorization and background colorization. In instance colorization, each object instance is colorized. While in background colorization, the remaining regions which do not belong to any object instance are colorized. This divide-and-conquer strategy leverages the natural advantages of the segmentation approach, making the colorization of a complex scene sketch resolvable.

Dataset. We use SketchyScene as the basic dataset to study the segmentation and colorization problem. SketchyScene contains 7,264 scene templates, each scene template corresponding to a color reference image. Moreover, SketchyScene provides the ground-truth for both semantic segmentation and instance segmentation of 7, 264 scene sketches. Figure 3 shows a typical example of a scene sketch in SketchyScene. See the supplementary material for more details of SketchyScene.

4 SEGMENTATION AND CAPTIONING

4.1 Instance Segmentation

Formulation. Instance segmentation segments individual object instances, possibly of the same object class in a scene sketch. This is challenging, especially when the instances of the same class have occlusions, *e.g.*, two trees growing together, as shown in Figure 3. Apart from using pixel-level class labels as supervision, spatial information like object bounding box is necessary. Generally, in an instance-level segmentation task, each unknown instance I in the input image can be denoted as a tuple $[B, L, M]$. Here, M is a binary mask covering the instance, L is a class label, and B is the 4-D vector encoding the position and size of the bounding box in the format of $[x, y, H, W]$, where $[x, y]$ indicates the top-left corner and H and W

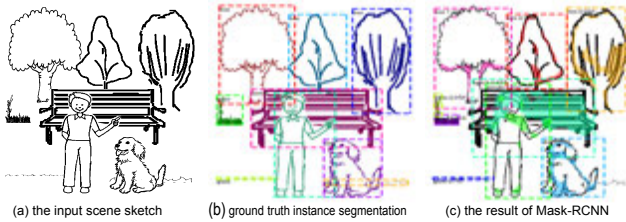(a) the input scene sketch   (b) ground truth instance segmentation   (c) the result of Mask-RCNN

Fig. 4. An illustration of instance segmentation using the state-of-the-art model Mask-RCNN [He et al. 2017] on a representative scene sketch of SketchyScene. Note that though the bounding boxes in (c) from the Mask-RCNN model are roughly aligned with the ground truth labels (b), the Mask-RCNN model misclassifies the strokes of the person as the bench class and fails to follow the black lines at the pixel level.

represent the height and width, respectively. Our goal of instance segmentation is to assign pixels of a scene sketch image to a specific instance mask M located in an inferred B.

Unlike natural images, a sketch only consists of black lines and a white background. Given that only black pixels convey semantic information, our problem for instance segmentation of a scene sketch is defined as predicting $[B, L, M]$ for each black pixel. Taking the rightmost image of Figure 3 as an example, when segmenting trees, duck, house and cloud, every black pixel should be assigned to a specific instance mask M with a class label L, located in B, while the remaining white pixels are treated as background.

Challenges. Segmenting a sketchy scene is challenging mainly due to the sparsity of visual features. First, a scene sketch image is dominated by white pixels. For the $7, 264$ examples of SketchyScene, the average background ratio is 87.83%. The remaining pixels belong to foreground classes. The classes, as measured by their sketch pixels, are thus quite imbalanced. Second, segmenting occluded objects in sketches is much harder than in natural images, where an object instance often contains uniform colors or texture so that context information can help in segmentation. Unfortunately, such cues do not exist in a scene sketch.

The sparsity of visual features causes the instance segmentation models, designed for natural images, perform poorly on scene sketches. In an initial experiment, we found that the segmentation results are unsatisfactory even with the state-of-the-art model Mask-RCNN [He et al. 2017] (see a representative result in Fig. 4 (3)). Although Mask-RCNN can successfully detect the bounding boxes of object instances in a scene sketch, the binary masks often fall out of the black lines.

Methodology. Our initial study in [Zou et al. 2018] has reported a significant finding on a semantic segmentation task. That is, the challenge for sketch scene segmentation is mainly caused by the large area of background. In a model tailored to this property the background pixels should not contribute to the loss during training. During the inference, background pixels are assigned to an arbitrary class label, and are filtered out by the drawing mask of the input sketch for the final output. Using this strategy, the adapted DeepLab-v2 model [Zou et al. 2018] can improve the semantic segmentation MIoU (Mean Intersection over Union) by more than 10%. On the test examples, this model achieved the best performance (63.1% on MIoU) over DeepLab-V3 models [Chen et al. 2017a], SegNet [Badrinarayanan et al. 2017], and FCN-based [Long et al. 2015] model.
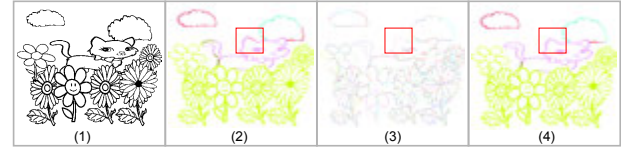

(1)   (2)   (3)   (4)

Fig. 5. Result of the adapted DeepLab-v2 model can be further refined by a postprocess based on edgelist, in which each entity is a skeleton of a drawing cluster. From left to right: (1) scene sketch, (2) output of the adapted DeepLab-v2 model in [Zou et al. 2018], (3) edgelist, and (4) post-processed result with edgelist.

Inspired by this finding, we have also tailored Mask-RCNN with this property, *i.e.*, ignoring the contribution of the background pixels when training the model. This adapted model also significantly improves the segmentation quality (Section 6).

Inspired by recent methods [Arnab and Torr 2017; Liang et al. 2015], which adopt a segmentation-first strategy, our final solution for instance segmentation is a framework which infers B and M based on semantic segmentation results. Specifically, we use Mask-RCNN to generate the bounding boxes B and L. We use the semantic segmentation results generated by an improved variation of the DeepLab-v2 model within the region of B to compute M.

Although the DeepLab-v2 model leverages a densely connected conditional random field (CRF) as the post-processor to refine the label inference, it still produces some discontinuous labels even on drawings which appear to belong to the same "edge" (see the ear of the cat in Figure 5). The degradation of CRF on sketches is mainly caused by the sparsity of sketches. CRF models the similarity of the class labels in an 8-connected neighborhood, while in a sketch a non-background pixel is usually affected by the neighboring pixels along sketch lines. This motivated us to use an assumption based on *edgelist* to further improve the results: the non-background pixels corresponding to the same drawing cluster (i.e., super-pixel consisting of only black pixels) should have the same class label.

An edgelist is defined as a set of single-pixel-width edges, in which each edge is a single-pixel-width skeleton of a drawing cluster as shown in Figure 5(3). Given a sketch, we first extract the single-pixel-width skeleton using the edge-link algorithm [Kovesi 2012]. After that, we compute the nearest skeleton (*i.e.*, an item of edgelist) for each non-background pixel and separate the sketch into a set of drawing clusters. We finally assign the majority class label of a drawing cluster to all the non-background pixels of the same cluster. We present the experimental results of this improved framework in Section 6.

4.2 Captioning

We develop a template-based algorithm for caption generation that starts from the output of the segmentation module (*i.e.*, a list of $[B, L, M]$, where B_i, L_i, and M_i contain the information of size, location, and class label of an object instance i). The steps of the algorithm are shown in Algorithm 1. It first classifies all the objects into three sets according to their class labels (we make this classification because SketchyScene was built by three sets of objects: objects related to weather or time, objects related to the site environment, and other objects related to the site environment). The first set **W** is for objects related to weather or time (*e.g.*, sunny or cloudy, night or
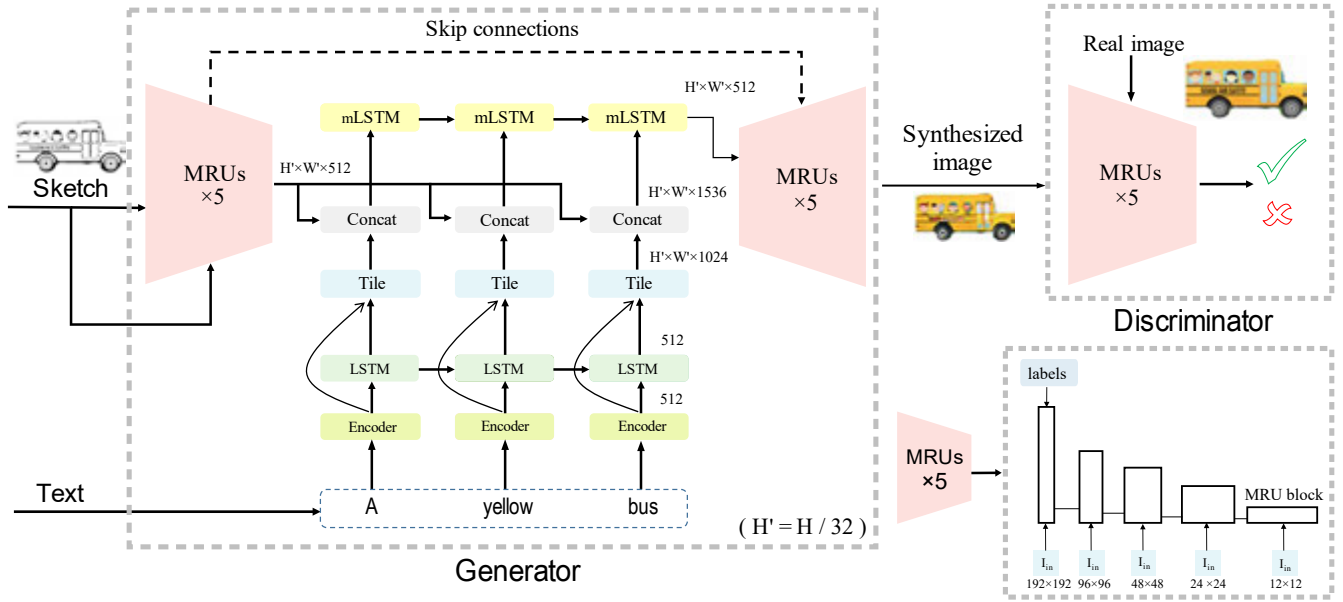
Fig. 6. Network architecture for object colorization, composed of a convolutional image encoder built on MRU blocks, a fusion module consisting of LSTM text encoders and multimodal LSTMs (mLSTM), a de-convolutional image decoder, and a MRU-blocks-based convolutional discriminator. The structure of the MRUs, which is shown at the right bottom corner, is inspired by SketchyGAN [Chen and Hays 2018].

day). The second set **E** is for objects which usually have a larger size and might determine the site environment of the scene sketch (*e.g.*, house, road, table, and bus). The last set is for the other objects **O**. Afterwards, we sequentially describe the objects from [**W, E, O**]. For objects from **W**, the algorithm generates a general weather description according to object labels, *e.g.*, "it is a sunny day". For objects from **E**, it produces the sentences which describe the information of class labels and absolute locations. A typical description is "There is a house in the center of the image". For objects from **O**, it produces the description of the formation of class labels and relative locations. A typical description is "A person is in front of the house".

Algorithm 1: Sketch Captioning

Input: pred_boxes B, pred_class_labels L, pred_class_masks M,
Output: caption T

1 *items* ←Node(B,L,M)
2 **for** *item* ∈ *items* **do**
3   **if** *item* ∈ **W** **then**
4        $T_{Weather}$←Weather(*item*)
5   **if** *item* ∈ **E** **then**
6        $T_{Environment}$←Environment(*item*)
7   **if** *item* ∈ **O** **then**
8   T_{Object}←Object(*item*)
9 T←$T_{weather}$+$T_{Envirment}$+T_{Object}
10 **return** T

The captioning module is applied to a human-machine interface as shown in the supplementary video. It is used to visualize the correspondence between an object in the input sketch and the corresponding sentence in the caption. Apart from captioning itself, there are two additional functions: 1) visualize if the sketch parsing results are accurate or not, and 2) visualize if the colorization goal is assigned to a desired object.

5 COLORIZATION

The colorization process contains two sequential steps: object instance colorization and background colorization. The object instance colorization step assigns target colors to the pixels belonging to segmented object instances, including their blank inner regions. Background colorization assigns colors to the remaining pixels.

5.1 Object Instance Colorization

Overview. The proposed framework for object instance colorization as shown in Figure 6 is a conditional GAN model consisting of a generator G and a discriminator D. G takes as input an object sketch image I and its corresponding language description $S = \{w_1, w_2, ..., w_t..., w_T\}$, where w_t are individual words in the sentence, and generates a color image. Compared with the generators in existing literature [Chen and Hays 2018; Isola et al. 2017], which learn mappings from an input sketch or image to an output image, our generator G models the interaction among the text description, visual information, and spatial relationships, and finally fuses the multi-modal features together. The generation of the color images is controlled by the text information. The discriminator, which is the opponent of the generator, is fed with both the generated images and the real color images at the same time, and serves the function of judging whether an image looks real or not.

Generator. The basic architecture of the generator is an encoder-decoder structure built on MRU blocks [Chen and Hays 2018]. It

consists of three modules: an image encoder which encodes the features of the $H \times W$ input sketch (segmented object sketch), a fusion module which fuses the text information into the image feature map generated by image encoder, and finally an image decoder which takes the fusion map produced by the fusion module and produces an $H \times W \times C$ map, where C is the number of color channels. The MRU block is first proposed in [Chen and Hays 2018], which uses a learned mask to selectively extract features from the input images. Its cascaded structure of MRU blocks allows the ConvNet to repeatedly and progressively retrieve the information from the input image on the computation path. In this work, we use MRU blocks in both the encoder and decoder. In our implementation, we use five cascaded MRUs to encode the $H \times W$ input sketch image into an $H' \times W'$ feature map ($H' = \frac{H}{32}, W' = \frac{W}{32}$) for the encoder. For the decoder, five symmetric MRUs are cascaded as a multi-layer de-convolutional network. Skip-connections are applied between the encoder and the decoder, concatenating the output feature maps from the encoder blocks to the output of the corresponding decoder blocks.

The fusion module fuses the text information in S into the $H' \times W'$ image feature map, and outputs an $H' \times W'$ fusion feature map. It is a core module of the generator and inserted into the bottleneck phase of the generator. The basic architecture of our fusion module is a convolutional multimodal LSTM, called recurrent multimodal interaction (RMI) model which was used to fuse the information of image referring expressions to segment out a referred region of the image [Liu et al. 2017a]. The typical characteristic of an RMI model is that the language model can access the image from the beginning of the language expression, allowing the modeling of the multimodal interaction. Our work uses RMI to mimic the human image colorization process. For each region in the input sketch, the fusion module reads the language feature map repeatedly until sufficient information is collected to colorize the target image region.

A concurrent system [Chen et al. 2017b] called LBIE (Language-Based Image Editing) has used a similar model to segment and colorize object parts of interest in an image conditioned by language descriptions. We have implemented the fusion module for our generator with both RMI and LBIE. We present the comparison results in the experimental section.

Discriminator. The discriminator D takes in a generated image and outputs the probability of the image being realistic. The structure of the discriminator follows SketchyGAN [Chen and Hays 2018] which uses four cascaded MRUs. It takes four types of scales of the real and generated images.

Loss and training. We use a hybrid loss following SketchyGAN [Chen and Hays 2018] which includes a GAN loss and an auxiliary classification loss in the Discriminator, a GAN loss, an auxiliary classification loss, an L1 loss, a perceptual loss, and a diversity loss in the Generator. The auxiliary classification loss can improve the ability of both the Discriminator and the Generator, and further enhances the quality of the synthesized images. The L1 distance loss is for comparison between the synthesized image and the ground-truth cartoon image. The perceptual loss and diversity loss are for generating diverse results.

5.2 Background Colorization

The background colorization network takes as input the result of object instance colorization, and infers the colors of the background regions to produce the final colorful image. The network architecture still employs a cGAN structure similar to that used for object instance colorization, with some modification as detailed below.

Generator. The architecture of the generator is similar to that shown in Figure 6. We replace MRU blocks with residual blocks [He et al. 2016]. Both the encoder and decoder use five cascaded residual blocks. The residual unit numbers of the five residual blocks for the encoder are {1, 3, 4, 6, 3} ({3, 6, 4, 3, 1} for the decoder). We make this change because an MRU block, which uses a binary mask to selectively extract features from a sketch, is more suitable for sketch images than color images. Our initial experiments confirm this speculation. The colorized backgrounds produced by the generator shown in Figure 6 (*i.e.*, the generator used for object instance) usually have sharp region boundaries, leading to relatively poor visual effects (also see the results shown in the second column from right of Figure15). Apart from the use of residual blocks, we have also conducted some experiments with the network replacing MRU blocks with the encoder blocks in [Isola et al. 2017].

Discriminator. For the discriminator, we use the architecture shown in Figure 7. It is a combination of the RMI model and five cascaded encode blocks used in pix2pix [Isola et al. 2017]. Unlike the generator, each of the five cascaded encoder blocks contains a single layer, which decreases the complexity of the entire cGAN. In addition, the input text information is also fused into the image feature by the RMI model. The joint modeling of the text and image helps the discriminator make a judgment monitored by the text information. Our experimental results (shown in supplementary materials) show that the joint modeling improves the capability of the whole network. Note that we do not use the RMI model to fuse the text information for the discriminator of object instance colorization. We simplify the discriminator because the discriminator with the fusion module does not improve the colorization significantly on single object sketches with lower resolution (192×192) in our initial experiments.

Loss and training. As for the loss and training, we follow the scheme of Pix2Pix [Isola et al. 2017]. We only use a conditional GAN loss as well as a L1 distance loss. The diversity loss is not used here as we do for object colorization, because we expect to suppress the diversity of the generated background.

6 SEGMENTATION EVALUATION

We conducted our experiments for segmentation on SketchyScene. The entire dataset, including 7,264 unique scene sketch templates, was randomly split into training (5,616), validation (535), and test (1,113) datasets.

Segmentation models. We compared four types of frameworks: Mask-RCNN (Model-1), Mask-RCNN + w/o BG (Model-2), Mask-RCNN + edgeList (Model-3), and Mask-RCNN + adapted DeepLab-v2 + edgeList (Model-4). These four frameworks are abbreviated as Model-1 to Model-4 below. Model-1 is extended from Faster-RCNN [Ren et al. 2015b] by adding a parallel object mask prediction branch and is one of the most advanced methods proposed for instance segmentation on natural images. Model-2 is an adapted
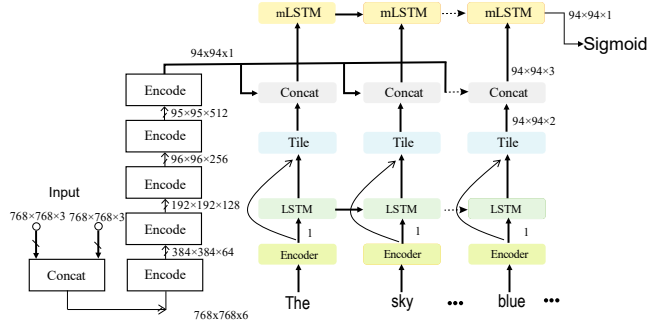
Fig. 7. Structure of the discriminator for background colorization. The structure combines five cascaded encode blocks used in the pix2pix network [Isola et al. 2017] and RMI model.
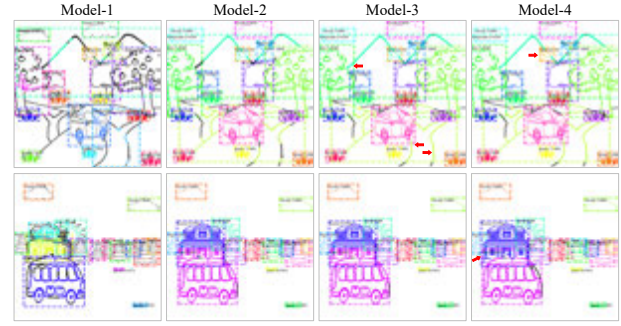

Fig. 8. Qualitative results of four competitors for instance segmentation. The improvement by ignoring background pixels is clear by comparing Model-1 and Model-2. In the third and fourth column, the benefits from edgelist and class labels can be found from the areas where red arrows point to.

Mask-RCNN model, where background pixels do not contribute to the loss during training. Model-3 is extended from Model-2 by adding a post-processing using edgelist. Model-4 is a combination model of the adapted DeepLab-v2 model [Zou et al. 2018] and Model-3, which means only the mask pixels, whose semantic label from adapted DeepLab-v2 is the same with the predicted label from Mask-RCNN, will be retained. The results of all the four models were filtered with the binary mask of non-background pixels. Similar to [He et al. 2017], we use *AP* (MS COCO metric), AP_{50} and AP_{75} (PASCAL VOC metrics) as evaluation metrics. It is worth mentioning that except for Model-1, all others: Model-2, Model-3, and Model-4 are our methods. The final segmentation model of LUCSS uses Model-3.

Implementation details. We implemented the proposed technique using Deeplab-v2 on Python 3 and TensorFlow, based on a ResNet101 backbone. The initial learning rate was set to 0.0001 and mini-batch size to 2. We set the maximum training iterations as 50K and the choose SGD (Stochastic Gradient Descent) as the optimizer. We resized the images as 768 × 768. The hyper parameters $\sigma_\alpha, \sigma_\beta$, and σ_γ of denseCRF in the adapted DeepLab-v2 were set to 7, 3, and 3, respectively.

Results. Table 1 shows the performance of the different models. Clearly, Model-2 performs much better than Model-1. It indicates that ignoring background pixels improves the segmentation task, which confirms the conclusion we made in [Zou et al. 2018]. Although Model-3 achieves slightly higher average quantitative values than Model-2, we find the improvement due to the use of edgelist usually occurs on some non-background pixels, which represent salient category characteristics such as the ear of a cat as shown in Figure 5. The improved performance is thus visually clearer. Model-4 obtains slightly better AP value in validation set compared to Model-3, but it gets worse in test set. Figure 8 shows some segmentation results of the above compared methods. More results can be found in the supplementary materials.

7 COLORIZATION EVALUATION

7.1 Main Results

We compared LUCSS to LBIE [Chen et al. 2017b] for scene sketch colorization, since to the best of our knowledge, LBIE is the only existing colorization work with the same goals as ours, i.e., taking a

Table 1. Instance segmentation mask AP on SketchyScene. Model-2, Model-3, and Model-4 are three models proposed by the work. *AP* (MS COCO metric), AP_{50} and AP_{75} (PASCAL VOC metrics) are the same metrics as those used in MaskRCNN.

Model	val			test		
	AP	AP_{50}	AP_{75}	AP	AP_{50}	AP_{75}
Model-1	10.03	29.20	4.86	12.99	36.76	5.93
Model-2	63.01	79.97	68.18	62.32	77.15	66.76
Model-3	63.78	**80.19**	68.88	**63.17**	**77.45**	**67.60**
Model-4	**64.38**	79.24	**69.10**	58.55	71.31	62.29

commented sketch as input, and outputting a color image. In general, the architecture of LBIE is similar to the RMI based architecture used for object colorization of LUCSS. Both of these two architectures are able to colorize an entire input sketch in a single step by implicitly inferring the correspondence between language descriptions and objects (or object parts). We therefore also include the RMI based architecture shown in Figure 6, which takes as input a scene sketch and a description of scene-level colorization requirement, and generates a color image in a singe step instead of two steps used by LUCSS, as another baseline method. In the following paragraphs of this section, these three comparison approaches are called LUCSS, LBIE, and sRMI for short.

7.1.1 Data Collection.

Data collection for LUCSS: object colorization. We collected three modalities of data to train the models of object colorization: color object instance, edge map, and text description (caption). We extracted color object instances and their edge maps from the 5,800 reference cartoon style images of SketchyScene. Figure 9 illustrates how the training data was prepared. We first leveraged Mask-RCNN trained on MS COCO to detect color object instances and then cut them out of the reference images. For each instance, we fused the filter responses of Hed [Xie and Tu 2015] and X-DoG [Winnemöller 2011] as the training edge map (object sketch). In total we collected

3,739 sets of object examples, each set consisting of a caption authored by crowd workers, a 192 × 192 color image, and a 192 × 192 edge map. The captions covered object instances from 20 categories, namely, moon, sun, cloud, house, bench, road, bus, car, bird, people, butterfly, cat, chicken, cow, dog, duck, sheep, tree, rabbit, and pig, which describes object instances in 15 different colors. The number of colors varies from one to three for each individual object instance (e.g., there are two colors for a red bus with gray windows). We further split the 3,739 sets of examples into two parts: 2,814 sets of examples used for training data, 357 sets for validation, and 568 sets for test.
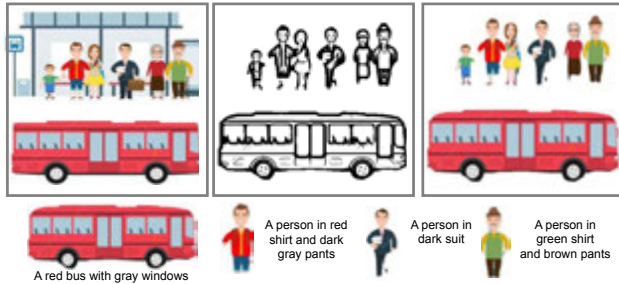

Fig. 9. Illustration of training data collection for object instance colorization. Top row: a reference image from SketchyScene (left); filter response of Hed [Xie and Tu 2015] and X-DoG [Winnemöller 2011] (middle); cut-out color object instances (right). Bottom row: user-written descriptions of the representative color object instances.

Data collection for LUCSS: background colorization. For background colorization, we used the following three modalities of training data: cartoon style color images, their captions, and color foreground object instances with blank background (See the left-most image of Figure 14 for an illustration). We used two types of strategies to extract the training data from the reference images of SketchyScene. The first one cut out objects from the reference images with mask-RCNN, and changed the color of all the pixels outside of the objects white. This strategy did not work well on some reference images where some object instances could not be detected. For these reference images, we used the other strategy: using the boundary of an object instance in the scene sketch as the boundary of the corresponding color object in the corresponding reference image (the object correspondence can be obtained from SketchyScene). We employed 5 workers to generate the captions. For each set of examples, the color details of two / three components, "sky", "land", and/or "background", were described. In total, we collected 1,328 sets of training examples. We used 1,200 sets of examples for training, the remaining 128 sets were equally split for validation and test. All the images were resized to 768 × 768 pixels.

Data collection for sRMI and LBIE. Similar to the data collection for LUCSS, we collected three types of data for sRMI and LBIE: (1) color images selected from the reference images of SketchyScene (we selected the reference images which were in the training set of SketchyScene and had the majority pixels belonging to the above-mentioned 20 categories), (2) corresponding edge maps, and (3) corresponding captions. To collect the captions in scales, we developed an on-line system to assist the crowd workers in generating captions for cartoon scenes. To ensure the description style is close to that

generated by the captioning algorithm, we used a template-based approach for caption generation. More specifically, each sentence of a caption comes from a repository of caption templates, which describes the instances and their corresponding quantities, colors, and spatial relationship as Algorithm 1. In this way, We employed 24 workers and collected 1,328 sets of examples, 1,200 of which were used for the training and the rest for validation.

Test data. We selected 100 scene sketches from the test set of SketchyScene as the test data (the other scene sketches do not depict reasonable scenes if the objects outside the selected 20 categories are removed from the scene sketch). Each scene sketch we produced had up to three captions (varying on different evaluation tasks). Apart from being used to evaluate the colorization performance of LUCSS, LBIE, and sRMI, this test data has also been used for the experiments on the analysis of the components of LUCSS.

7.1.2 Experimental Settings. For object instance colorization, we used a batch size of 2 and trained with $100K$ iterations. In the initial phrase of the training, we used the ADAM optimizer [Kingma and Ba 2014] and set the learning rate of generator at 0.0002 and that of discriminator at 0.0001. After $50K$ iterations, we adjusted the learning rate of discriminator to 0.0002. For background colorization, we used batch size of 1 and trained with $100K$ iterations. We set the initial learning rate for both the generator and discriminator at 0.0002 and reduced it by 75% after each $20K$ iterations. We set the iteration number of LSTM at 15, the cell size of mLSTM at 512 for both object instance and background colorization modules.

sRMI and LBIE. We trained $100K$ iterations for both sRMI and LBIE. We set the initial learning rate of both generator and discriminator at 0.0002 and reduce it by 75% after each $20K$ iterations for sRMI. We set the iteration number of LSTM at 200, the cell size of mLSTM at 128 for sRMI. The iteration numbers of LSTM in LBIE were adapted to the input captions. We set other experimental settings of LBIE by following [Chen et al. 2017b].

7.1.3 Comparison Results. We conducted two sets of experiments to evaluate the performances of LUCSS, sRMI, and LBIE.

Single caption. In the first set of experiments, we compared LUCSS, LBIE, and sRMI on the 100 scene sketches of the whole test data. For each test scene sketch, we produced one caption (by re-editing the results of captioning). Figure 10 shows the visual results of these three competitors on some representative test examples.

Generally, LUCSS outperformed both LBIE and sRMI overwhelmingly. LUCSS colorized both objects (e.g., the roads, houses, clouds, moons, trees in the results), and backgrounds (the green land and sky) with smooth colors, while LBIE and sRMI generated a lot of color regions which covered more than one objects. sRMI relatively performed better than LBIE in some cases. We take the results in the third row of Figure 10 for example. sRMI colorized the sky and grass with smooth blue and green respectively, while the color generated by LBIE was mixed with multiple colors (e.g. white, yellow and pink).

With respect to *faithfulness*, which measures whether the colorization follows the language instructions, LUCSS achieved significantly better performance than LBIE and sRMI as well. We take the example in the fourth row of Figure 10 to explain our observation. LUCSS

colorized all the clouds light gray, while LBIE colorized them into different colors including white and green, and sRMI colorized them into white, green and orange; LUCSS painted the lady red following the caption, while LBIE painted her into blue and red gradient, and sRMI colorized her with red hair and orange skirt. Relatively, sRMI performed a little better than LBIE on some examples. For example, sRMI can colorize the sky and grass in the third row of Figure 10 correctly, while LBIE failed to meet the demand of the caption.

The poor results of LBIE and sRMI may be mainly due to two factors. First, both LBIE and sRMI did not achieve accurate instance segmentation, which can be indicated by the tree at the most right side (LBIE: white and red; sRMI: green and blue), and the person (LBIE: blue skirt and red legs; sRMI: red hair, orange skirt and green legs). In contrast, LUCSS achieved accurate instance segmentation, since the boundaries of the resulting color objects were consistent with those in the sketch. Second, either LBIE or sRMI did not achieve accurate correspondence between the description and the objects. We can see the evidence from the sun (LBIE: blue sun; sRMI: white). In LUCSS, the correspondence is naturally obtained by the captioning module.

Multiple captions. In the second set of experiment, we randomly selected 30 examples from the test dataset and produced three captions for each scene sketch. We then evaluated LUCSS, LBIE, and sRMI on all the three captions for each scene sketch. Some representative results are shown in Figure 11. The results of LUCSS generally met the user-specified colorization requirements. The results of LBIE on three different captions almost stayed the same. Although sRMI can response to the changed captions in some cases, it colored most objects in wrong colors. It indicate that both LBIE and sRMI failed to obtain the correspondence between the captions and objects. In [Chen et al. 2017b], LBIE showed the ability to learn the correspondence between words and objects, which is possibly because the captions in [Chen et al. 2017b] contain much shorter sentences than these in this study (captions in our study typically contain more than 6 sentences, much longer than two or three sentences in [Chen et al. 2017b]).

7.2 Ablation Experiments

In this section, we design various experiments to analyze the two major components of LUCSS.

7.2.1 Object Colorization.

LBIE versus MRU-RMI. In this set of experiment, we evaluated the performance of two types of models, LBIE and MRU-RMI, for object colorization. The test data contained 568 sets of examples as discussed in the data collection section above. Each set of test examples contained an object sketch and a corresponding caption. LBIE is the same network as that used in the previous experiments. MRU-RMI is also the same network as sRMI, which was used to colorize a whole scene sketch in the previous experiments.

Figure 12 shows some representative results. It can be seen that both LBIE and MRU-RMI can learn an implicit part-level segmentation and colorize the object parts with the colors instructed by the captions (e.g., both LBIE and MRU-RMI assigned blue to the windows of the bus). Relatively, MRU-RMI achieved better performance on inferring the correspondence between words in the caption and
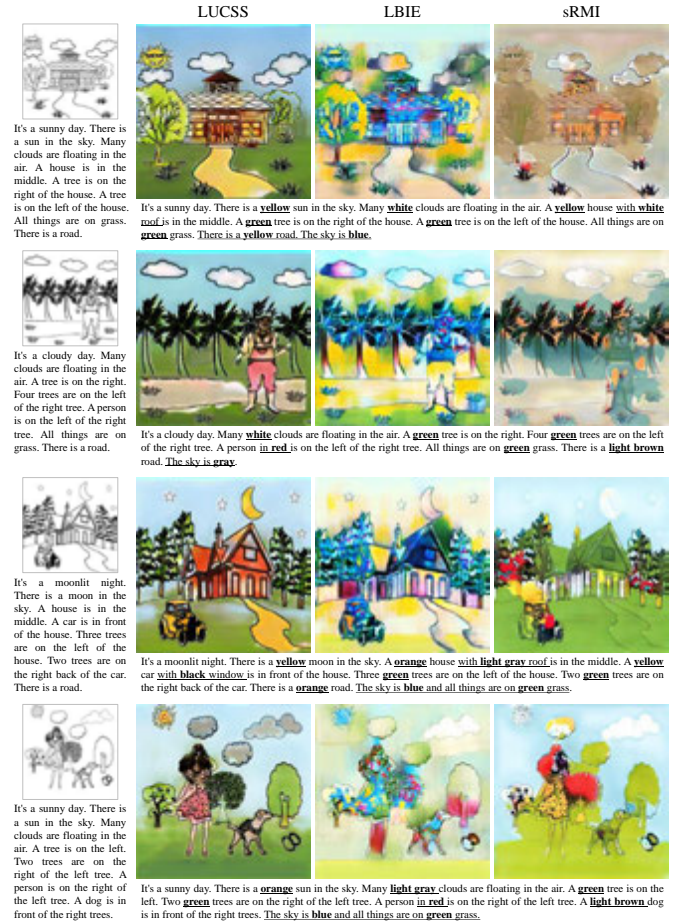
Fig. 10. LUCSS *vs.* LBIE *vs.* sRMI on various inputs. An input scene sketch and the automatically generated caption are shown on the left for each set of example. User edited captions are shown at the bottom of each colorization result. Each row shows the results produced by the three competitors. More results can be found in supplementary materials.

object parts. This can be told from the colorization results for the roof of the house and the windows of the bus. Combining the results in both Figure 12 and 11, we can see that MRU-RMI is powerful for object-level sketch images and its performance would be significantly degraded when the sketch image size goes up to a large-size complex scene sketch. Moreover, for the face region of the person in Figure 12, we can see that MRU-based network also outperformed the atrous-convolution-based LBIE on the encoding of object-level sketch image features (performance of MRU based network is significantly degraded when the complexity and the size of sketch image increase).

MRU versus ResNet blocks versus Pix2Pix. In this set of experiments, we evaluated three different types of backbones for the encoder and decoder of the architecture used for object colorization. The first type of backbone is MRU, which is used by SkecthyGAN[Chen and Hays 2018] as well as the encoder and decoder for object colorization of LUCSS. ResNet denotes the residual block
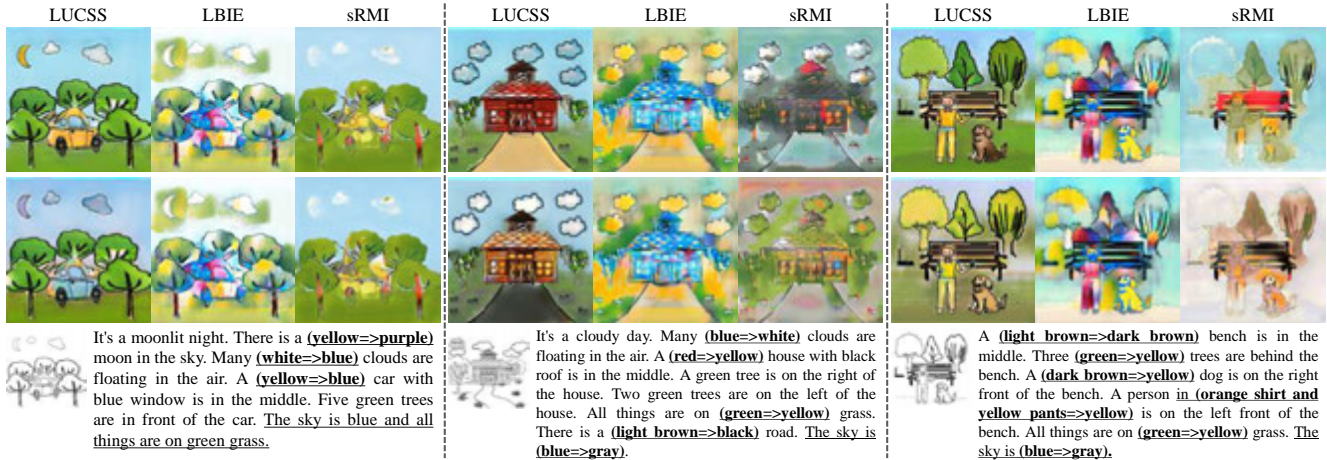
It's a moonlit night. There is a **(yellow=>purple)** moon in the sky. Many **(white=>blue)** clouds are floating in the air. A **(yellow=>blue)** car with blue window is in the middle. Five green trees are in front of the car. <u>The sky is blue and all things are on green grass.</u>

It's a cloudy day. Many **(blue=>white)** clouds are floating in the air. A **(red=>yellow)** house with black roof is in the middle. A green tree is on the right of the house. Two green trees are on the left of the house. All things are on **(green=>yellow)** grass. There is a **(light brown=>black)** road. <u>The sky is **(blue=>gray)**.</u>

A **(light brown=>dark brown)** bench is in the middle. Three **(green=>yellow)** trees are behind the bench. A **(dark brown=>yellow)** dog is on the right front of the bench. A person in **(orange shirt and yellow pants=>yellow)** is on the left front of the bench. All things are on **(green=>yellow)** grass. <u>The sky is **(blue=>gray)**.</u>

Fig. 11. Colorization results by LUCSS, LBIE, and sRMI with two different sets of user-specified colorization requirements for each scene sketch. For each set of comparison, an input scene sketch is at the bottom, with the corresponding caption. Text without underline is automatically generated, while underlined text is added by users. Arrows in the parentheses represent the change of colorization requirements (text on the left side of the arrows correspond to the results in the top row; text on the right side of the arrows correspond to the results in the bottom row).
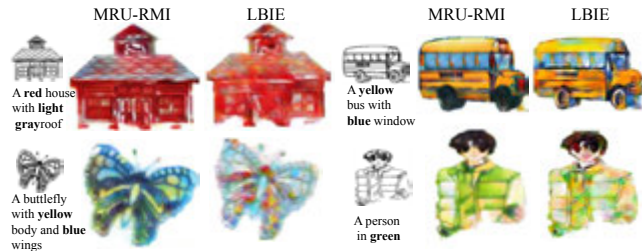

Fig. 12. Network architecture evaluation for object colorization: MRU-RMI versus LBIE. MRU-RMI (shown in Figure 6) denotes the architecture used by LUCSS for object colorization. It uses MRU backbone based encoder & decoder and RMI based fusion model. LBIE is the model used in [Chen et al. 2017b]. Input scene sketches and language descriptions are shown on the left of each example.

proposed by [He et al. 2016]. Pix2Pix denotes the convolution block used by the encoder and decoder of [Isola et al. 2017]. For comparison purposes, we replaced the cascade blocks in Figure 6 with ResNet/Pix2Pix, and produced results on all the 568 sets of test examples.
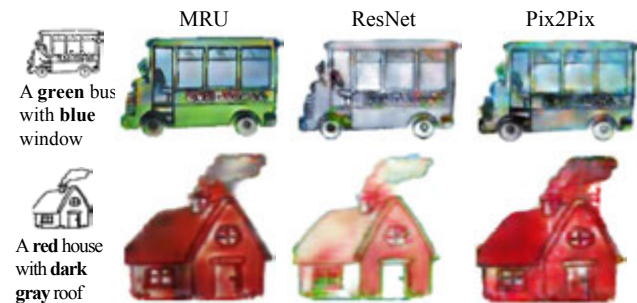

Fig. 13. Network backbone evaluation for object colorization : MRU versus ResNet versus Pix2Pix. MRU based image encoder & decoder outperforms both ResNet and Pix2Pix based image encoder & decoder.

In Figure 13, we show the results of two representative examples. MRU achieved better performance than both Pix2Pix and ResNet. Specifically, MRU generated clearer texture and object part boundaries compared to ResNet and Pix2Pix (e.g., see the window and the body of the colorized bus). In the aspect of following language instruction, the results of MRU are also superior to those of ResNet and Pix2Pix. This also indicates that the whole network based on MRU can infer more accurate correspondence between the caption and image content using the image features extracted by MRU.

7.2.2 Background Colorization. In this section, we first investigate the architecture of the GAN for the background colorization task, and then study what kind of backbone is more appropriate for this task.

***LBIE* versus *ResNet-RMI*.** In this set of experiments, we compared two types of architectures, which can be used for background colorization. The first architecture is LBIE. The other is the architecture discussed in Section 5.2. We call this architecture ResNet-RMI in this section since the backbone of the generator uses ResNet blocks. Both LBIE and ResNet-RMI were trained with 1,200 sketches in the training dataset, and tested on the test dataset.

In Figure 14, we illustrate the results of LBIE and ResNet-RMI on two sets of representative examples. Both LBIE and ResNet-RMI successfully colorized the regions of sky and grass with the colors specified in the captions. This indicates that the encoding modules of image, text and feature fusion modules of LBIE and ResNet-RMI work well. However, the results by ResNet-RMI were visually more pleasing than LBIE. It may be caused by the fact that LBIE uses an asymmetrical encoder-decoder architecture (the encoder of LBIE uses atrous convolution while the decoder uses regular convolution).

***MRU* versus *ResNet* versus *Pix2Pix*.** In this set of experiments, we study which backbone network is more suitable for large-scale image background colorization. We still used the three types of backbone network: MRU, ResNet, and Pix2Pix. Results on representative
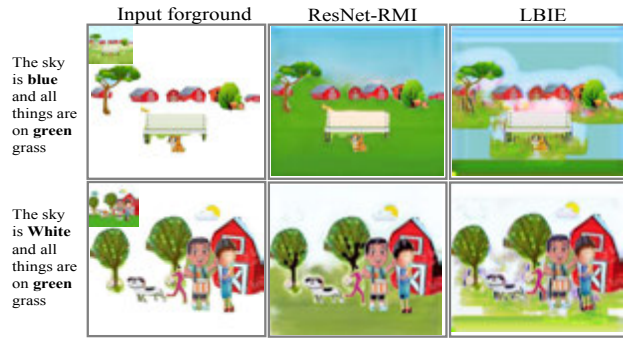
Fig. 14. Network architecture evaluation for background colorization : ResNet-RMI versus LBIE. It took as input the language descriptions on the left, and images containing color foreground objects and blank backgrounds. The reference images from which color foreground objects were cut out are shown on the top left corner of each input image.

examples are shown in Figure 15. We can see ResNet outperforms Pix2Pix in terms of visual effects, which can be explained by the fact that ResNet is much deeper than Pix2Pix (a ResNet block has three convolution layers while a Pix2Pix block only has a single convolution layer). ResNet also has visually better results than MRU (e.g., MRU generated some artifacts surrounding foreground objects, while ResNet doesn't), mainly because MRU is specialized for object-level sketches. When the MRU based network is used for a large size of color image, its performance is degraded (we can get the evidence by comparing the results in Figure 11 and Figure 12, additional evidence can be seen in SketchyGAN[Chen and Hays 2018]).
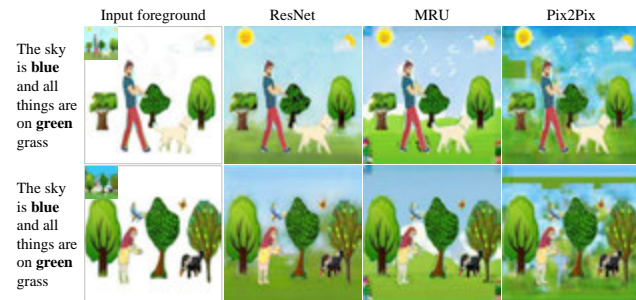

Fig. 15. Network backbone evaluation for background colorization : ResNet *vs.* MRU *vs.* Pix2Pix. ResNet outperforms both MRU and Pix2Pix when colorizing blank backgrounds surrounding color foreground objects.

7.3 Human Evaluation

To further justify the captioning and colorization results of LUCSS, we have designed and carried out two sets of user experiments: *faithfulness study* and *effectiveness study*. We recruited 11 participants for the studies. All participants are undergraduates with no prior knowledge of this project.

Each set of experiments covers the two sub-steps of LUCSS: object colorization and background colorization. We did not conduct a study for system level performance of LUCSS, LBIE, and sRMI (see Section 7.1) because LUCSS obviously outperforms LBIE and LBIE outperforms sRMI on all the examples of the test dataset.
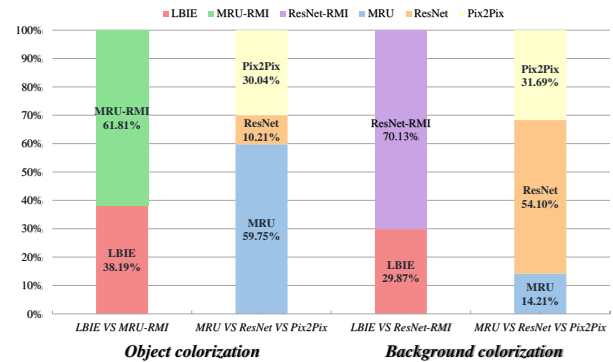

Fig. 16. Quantitative results for the faithfulness study. LUCSS uses the RMI model in general, MRU backbone for object colorization, and ResNet blocks for background colorization.

In the object colorization task, we randomly select 60 sketches from 20 classes as the input (3 sketches per class) and compare the results using two models: LBIE and MRU-RMI (our model). In the background colorization task, we random select 20 scene sketches as the input and compare the results via two models: LBIE and ResNet-RMI (our model). In both tasks, we further the performance of three backbones: MRU, ResNet blocks, and Pix2Pix.

In the *faithfulness study*, our goal is to evaluate whether the colorization results are consistent with the text description. Each participant was given the input caption with the corresponding colorization results from one of the two models, or one of the three backbones. We asked the participant to pick out the colorization result which best fits the text description.

In the *effectiveness study*, we compared the overall visual quality of the colorization results. Each participant was given the original sketch with the corresponding colorization results from one of the two models, or one of the three backbones. We asked the participants to pick out the most visually pleasing color image for the sketch.

Overall, in each study, we have collected $11 \times 20 \times 3 = 660$ trials for the object colorization task and $11 \times 20 = 220$ trials for the background colorization task. For both studies, we compare the selection rate of different models and backbones in Fig. 16 and 17. For both tasks, our MRU-RMI model and ResNet-RMI model greatly outperform the LBIE model in user evaluation. Based on our user evaluation, LUCSS is more faithful and produces more visually pleasing results to users. For object colorization, the MRU backbone stands out; while for background colorization, the ResNet backbone outperforms the rest.

8 CONCLUSION, DISCUSSION, AND FUTURE WORK

Understanding low-level scene sketches containing multiple sketched objects is a rarely studied problem. This problem is very challenging, especially when objects occlude each other in the depicted scene. The sparse nature of scene sketches leads to inferior performance of current models, even advanced ones, that are designed for natural images. In this work, we proposed a system called LUCSS to study how we can make machines understand complex scene sketches by adapting the existing powerful deep models for natural images to
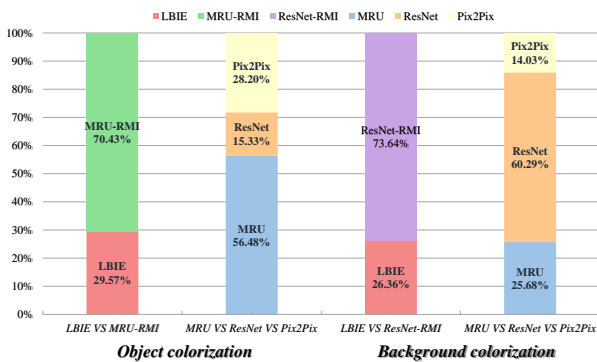
Fig. 17. Quantitative results for the effectiveness study. LUCSS uses the RMI model in general, MRU backbone for object colorization, and ResNet blocks for background colorization.

sketches, as well as how this understanding can benefit related applications. We have focused language-based interactive colorization of scene sketches.

The current performance of LUCSS is limited by the segmentation performance since the highest mask AP is near 60%. It is necessary to propose more powerful algorithms by incorporating the characteristics of scene sketches with advanced models. Another limitation is that our current human-computer interaction needs improvement, though it is currently a practical solution considering the challenges of instance segmentation and inferring accurate correspondence between language descriptions and objects in the scene. As these challenging problems are solved by new datasets and more powerful models, LUCSS has the potential to become a far more mature system that could potentially generate fine-grained sketch colorization like generating a boy with ruddy cheek, or animating objects in a scene sketch by human speech commands.

Image synthesis often has problems when data is left out. Since LUCSS is limited by training data, LUCSS generates cartoon images rather than natural images. Complex sentences are input to LUCSS which then processes them, filling in any missing detail such as omitted colors, and outputs an image.

This raises a question: is it possible to synthesize user-customized scene-level natural images by adding user-drawn sketches to the images of MS COCO's dataset to corroborate the existing language descriptions?

LUCSS has great application prospects in the field of child education, accessibility, and media production. For instance, by translating the auto-generated caption text to human voice, LUCSS has the potential to read scene sketches to children and the blind. Via the interactive colorization system, both children and adults could create their own cartoon story books. In addition, the involved techniques in LUCSS may be useful in CAD and Virtual Reality (VR) industries. With text (voice) commands, LUCSS unlocks the potential to easily change the color schemes of a sketch scene and virtual environments.

REFERENCES

Xiaobo An and Fabio Pellacini. 2010. User-Controllable Color Transfer. In *Computer Graphics Forum*, Vol. 29. Wiley Online Library, 263–271.

Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. In *CVPR*.

Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning*, Vol. 70. 214–223.

Anurag Arnab and Philip H. S. Torr. 2017. Pixelwise Instance Segmentation with a Dynamically Instantiated Network. In *CVPR*. 879–888.

Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. 2017. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* 39, 12 (2017), 2481–2495.

Raffaella Bernardi, Ruket Cakici, Desmond Elliott, Aykut Erdem, Erkut Erdem, Nazli Ikizler-Cinbis, Frank Keller, Adrian Muscat, and Barbara Plank. 2016. Automatic Description Generation from Images: A Survey of Models, Datasets, and Evaluation Measures. *arXiv preprint arXiv:1601.03896* (2016).

Huiwen Chang, Ohad Fried, Yiming Liu, Stephen DiVerdi, and Adam Finkelstein. 2015. Palette-based photo recoloring. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 139.

Guillaume Charpiat, Matthias Hofmann, and Bernhard Schölkopf. 2008. Automatic Image Colorization Via Multi-modal Predictions. In *Proceedings of the 10th European Conference on Computer Vision: Part III (ECCV)*. 126–139.

Hwann-Tzong Chen, Huang-Wei Chang, and Tyng-Luh Liu. 2005. Local discriminant embedding and its variants. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, Vol. 2. IEEE, 846–853.

Jianbo Chen, Yelong Shen, Jianfeng Gao, Jingjing Liu, and Xiaodong Liu. 2017b. Language Based Image Editing with Recurrent Attentive Models. *CoRR* abs/1711.06288 (2017).

Liang-Chieh Chen, Alexander Hermans, George Papandreou, Florian Schroff, Peng Wang, and Hartwig Adam. 2017a. MaskLab: Instance Segmentation by Refining Object Detection with Semantic and Direction Features. *CoRR* abs/1712.04837 (2017).

Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. 2016. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *arXiv:1606.00915* (2016).

Tao Chen, Ming-Ming Cheng, Ping Tan, Ariel Shamir, and Shi-Min Hu. 2009. Sketch2Photo: internet image montage. *ACM Trans. Graph.* 28, 5 (2009), 124:1–124:10.

Wengling Chen and James Hays. 2018. SketchyGAN: Towards Diverse and Realistic Sketch to Image Synthesis. *CoRR* abs/1801.02753 (2018).

Xinlei Chen and C. Lawrence Zitnick. 2014. Learning a Recurrent Visual Representation for Image Caption Generation. *CoRR* abs/1411.5654 (2014).

Z. Cheng, Q. Yang, and B. Sheng. 2015. Deep Colorization. In *ICCV*. 415–423.

Abhishek Das, Harsh Agrawal, Larry Zitnick, Devi Parikh, and Dhruv Batra. 2017. Human Attention in Visual Question Answering: Do Humans and Deep Networks Look at the Same Regions? *Computer Vision and Image Understanding* 163 (2017), 90–100.

Jia Deng, Wei Dong, Richard Socher, LiJia Li, Kai Li, and FeiFei Li. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR*.

A. Deshpande, J. Rock, and D. Forsyth. 2015. Learning Large-Scale Automatic Image Colorization. In *ICCV*. 567–575.

Jeff Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. 2014. Long-term Recurrent Convolutional Networks for Visual Recognition and Description. *CoRR* abs/1411.4389 (2014).

Mathias Eitz, James Hays, and Marc Alexa. 2012a. How do humans sketch objects? *ACM Trans. Graph.* 31, 4 (2012), 44–1.

Mathias Eitz, James Hays, and Marc Alexa. 2012b. How Do Humans Sketch Objects? *ACM Trans. Graph. (Proc. SIGGRAPH)* 31, 4 (2012), 44:1–44:10.

Kevin Frans. 2017. Outline Colorization through Tandem Adversarial Networks. *arXiv preprint arXiv:1704.08834* (2017).

Ross B. Girshick. 2015. Fast R-CNN. *CoRR* abs/1504.08083 (2015).

Yagmur Güçlütürk, Umut Güçlü, Rob van Lier, and Marcel A. J. van Gerven. 2016. Convolutional Sketch Inversion. *CoRR* abs/1606.03073 (2016).

David Ha and Douglas Eck. 2017. A neural representation of sketch drawings. *arXiv preprint arXiv:1704.03477* (2017).

Bharath Hariharan, Pablo Andres Arbelaez, Ross B. Girshick, and Jitendra Malik. 2014. Simultaneous Detection and Segmentation. In *ECCV*. 297–312.

Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross B. Girshick. 2017. Mask R-CNN. In *ICCV*. 2980–2988.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*. 770–778.

Seunghoon Hong, Dingdong Yang, Jongwook Choi, and Honglak Lee. 2018. Inferring Semantic Layout for Hierarchical Text-to-Image Synthesis. *CoRR* abs/1801.05091 (2018).

YiChin Huang, YiShin Tung, JunCheng Chen, SungWen Wang, and JaLing Wu. 2005. An Adaptive Edge Detection Based Colorization Algorithm and Its Applications. In *Proceedings of the 13th Annual ACM International Conference on Multimedia (MULTIMEDIA '05)*. 351–354.

Zhe Huang, Hongbo Fu, and Rynson W. H. Lau. 2014. Data-driven segmentation and labeling of freehand sketches. *ACM Trans. Graph.* 33, 6 (2014), 175:1–175:10.

Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. 2016. Let There Be Color!: Joint End-to-end Learning of Global and Local Image Priors for Automatic Image Colorization with Simultaneous Classification. *ACM Trans. Graph.* 35, 4 (2016).

Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. *CVPR* (2017).

Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980 (2014).

P. Kovesi. 2012. MATLAB and Octave Functions for Computer Vision and Image Processing. (2012). Available from: <http://www.peterkovesi.com/matlabfns/>.

Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. 2016. Learning Representations for Automatic Colorization. *CoRR* abs/1603.06668 (2016). http://arxiv.org/abs/1603.06668

C. Ledig, L. Theis, F. HuszÃąr, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. 2017. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Anat Levin, Dani Lischinski, and Yair Weiss. 2004. Colorization Using Optimization. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 689–694.

Xiaodan Liang, Yunchao Wei, Xiaohui Shen, Jianchao Yang, Liang Lin, and Shuicheng Yan. 2015. Proposal-free Network for Instance-level Object Segmentation. *CoRR* abs/1509.02636 (2015).

Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. *CoRR* abs/1405.0312 (2014).

Chenxi Liu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, and Alan L. Yuille. 2017a. Recurrent Multimodal Interaction for Referring Image Segmentation. *CoRR* abs/1703.07939 (2017).

Yifan Liu, Zengchang Qin, Zhenbo Luo, and Hua Wang. 2017b. Auto-painter: Cartoon Image Generation from Sketch by Using Conditional Generative Adversarial Networks. *CoRR* abs/1705.01908 (2017).

Yifan Liu, Zengchang Qin, Zhenbo Luo, and Hua Wang. 2017c. Auto-painter: Cartoon Image Generation from Sketch by Using Conditional Generative Adversarial Networks. (05 2017).

Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *CVPR*. 3431–3440.

Qing Luan, Fang Wen, Daniel Cohen-Or, Lin Liang, Ying-Qing Xu, and Heung-Yeung Shum. 2007. Natural image colorization. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*. Eurographics Association, 309–320.

Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan Yuille. 2015. Deep Captioning with Multimodal Recurrent Neural Networks (m-RNN). *ICLR* (2015).

Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei Efros. 2016. Context Encoders: Feature Learning by Inpainting. In *CVPR*.

Yingge Qu, Tien-Tsin Wong, and Pheng-Ann Heng. 2006a. Manga Colorization. *ACM Trans. Graph.* 25, 3 (2006).

Yingge Qu, Tien-Tsin Wong, and Pheng-Ann Heng. 2006b. Manga colorization. In *ACM Transactions on Graphics (TOG)*, Vol. 25. ACM, 1214–1220.

Alec Radford, Luke Metz, and Soumith Chintala. 2015. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *CoRR* abs/1511.06434 (2015).

Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. 2016. Generative Adversarial Text to Image Synthesis. In *Proceedings of The 33rd International Conference on Machine Learning*, Vol. 48. 1060–1069.

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015a. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems (NIPS)*.

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015b. Faster R-CNN: Towards Real-time Object Detection with Region Proposal Networks. In *NIPS*. 91–99.

Patsorn Sangkloy, Jingwan Lu, Chen Fang, Fisher Yu, and James Hays. 2017a. Scribbler: Controlling Deep Image Synthesis with Sketch and Color. In *CVPR*. 6836–6845.

Patsorn Sangkloy, Jingwan Lu, Chen Fang, Fisher Yu, and James Hays. 2017b. Scribbler: Controlling deep image synthesis with sketch and color. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Vol. 2.

Rosália G. Schneider and Tinne Tuytelaars. 2014. Sketch Classification and Classification-driven Analysis Using Fisher Vectors. *ACM Trans. Graph.* 33, 6 (Nov. 2014), 174:1–174:9.

Rosalia G. Schneider and Tinne Tuytelaars. 2016. Example-Based Sketch Segmentation and Labeling Using CRFs. *ACM Trans. Graph.* 35, 5 (2016), 151:1–151:9.

Evan Shelhamer, Jonathan Long, and Trevor Darrell. 2016. Fully Convolutional Networks for Semantic Segmentation. *CoRR* abs/1605.06211 (2016).

Casper Kaae Sønderby, Jose Caballero, Lucas Theis, Wenzhe Shi, and Ferenc Huszár. 2016. Amortised MAP Inference for Image Super-resolution. *CoRR* abs/1610.04490 (2016).

Zhenbang Sun, Changhu Wang, Liqing Zhang, and Lei Zhang. 2012a. Free Hand-Drawn Sketch Segmentation. In *ECCV*. 626–639.

Zhenbang Sun, Changhu Wang, Liqing Zhang, and Lei Zhang. 2012b. Sketch2Tag: automatic hand-drawn sketch recognition. In *ACM Multimedia*. 1255–1256.

Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. 2016. Pixel Recurrent Neural Networks. *CoRR* abs/1601.06759 (2016).

Domonkos Varga, Csaba Attila Szabo, and Tamas Sziranyi. 2017. Automatic Cartoon Colorization Based on Convolutional Neural Network. In *CMBI*.

Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2016. Show and Tell: Lessons learned from the 2015 MSCOCO Image Captioning Challenge. *CoRR* abs/1609.06647 (2016).

Changhu Wang, Jun Zhang, Bruce Yang, and Lei Zhang. 2011. Sketch2Cartoon: Composing Cartoon Images by Sketching. In *ACM MultiMedia*. 789–790.

Fang Wang, Le Kang, and Yi Li. 2015. Sketch-based 3D shape retrieval using Convolutional Neural Networks. In *CVPR*. 1875–1883.

Tomihisa Welsh, Michael Ashikhmin, and Klaus Mueller. 2002. Transferring Color to Greyscale Images. *ACM Trans. Graph.* 21, 3 (July 2002), 277–280.

Holger Winnemöller. 2011. XDoG: Advanced Image Stylization with eXtended Difference-of-Gaussians. In *Eurographics Symposium on Non-Photorealistic Animation and Rendering (NPAR '11)*. 147–156.

Wenqi Xian, Patsorn Sangkloy, Jingwan Lu, Chen Fang, Fisher Yu, and James Hays. 2017. TextureGAN: Controlling Deep Image Synthesis with Texture Patches. *CoRR* abs/1706.02823 (2017).

S. Xie and Z. Tu. 2015. Holistically-Nested Edge Detection. In *ICCV*. 1395–1403.

Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. *CoRR* abs/1502.03044 (2015).

Kun Xu, Kang Chen, Hongbo Fu, Wei-Lun Sun, and Shi-Min Hu. 2013. Sketch2Scene: sketch-based co-retrieval and co-placement of 3D models. *ACM Trans. Graph.* 32, 4 (2013), 123:1–123:15.

Kun Xu, Yong Li, Tao Ju, Shi-Min Hu, and Tian-Qiang Liu. 2009. Efficient affinity-based edit propagation using kd tree. In *ACM Transactions on Graphics (TOG)*, Vol. 28. ACM, 118.

Zhicheng Yan, Hao Zhang, Baoyuan Wang, Sylvain Paris, and Yizhou Yu. 2016. Automatic Photo Adjustment Using Deep Neural Networks. *ACM Trans. Graph.* 35, 2 (2016).

Chen Yao, Xiaokang Yang, Jia Wang, Song Li, and Guangtao Zhai. 2010. Patch-driven colorization. *Optical Engineering* 49, 1 (2010), 017001.

L. Yatziv and G. Sapiro. 2006. Fast image and video colorization using chrominance blending. *IEEE Transactions on Image Processing* 15, 5 (2006), 1120–1129.

Raymond A. Yeh, Chen Chen, Teck-Yian Lim, Mark Hasegawa-Johnson, and Minh N. Do. 2016. Semantic Image Inpainting with Perceptual and Contextual Losses. *CoRR* abs/1607.07539 (2016).

Taizan Yonetsuji. 2017. Paints Chainer. https://github.com/pfnet/PaintsChainer. (2017).

Qian Yu, Yongxin Yang, Feng Liu, Yi-Zhe Song, Tao Xiang, and Timothy M. Hospedales. 2017. Sketch-a-Net: A Deep Neural Network that Beats Humans. *International Journal of Computer Vision* 122, 3 (2017), 411–425.

Han Zhang, Tao Xu, and Hongsheng Li. 2017. StackGAN: Text to Photo-Realistic Image Synthesis with Stacked Generative Adversarial Networks. In *ICCV*. 5908–5916.

Jianhui Zhang, Yilan Chen, Lei Li, Hongbo Fu, and Chiew-Lan Tai. 2018. Context-based Sketch Classification. In *Expressive 2018*.

Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. 2016. Pyramid Scene Parsing Network. *CoRR* abs/1612.01105 (2016).

Luowei Zhou, Chenliang Xu, Parker Koch, and Jason J Corso. 2016. Image Caption Generation with Text-Conditional Semantic Attention. *arXiv preprint arXiv:1606.04621* (2016).

Changqing Zou, Qian Yu, Ruofei Du, Haoran Mo, Yi-Zhe Song, Tao Xiang, Chengying Gao, Baoquan Chen, and Hao Zhang. 2018. SketchyScene: Richly-Annotated Scene Sketches. In *ECCV*.