Pixelated Semantic Colorization

Jiaojiao Zhao¹ · Jungong Han² · Ling Shao³ · Cees G. M. Snoek¹

Received: 20 January 2019 / Accepted: 25 November 2019 / Published online: 7 December 2019
© The Author(s) 2019

Abstract

While many image colorization algorithms have recently shown the capability of producing plausible color versions from gray-scale photographs, they still suffer from limited semantic understanding. To address this shortcoming, we propose to exploit pixelated object semantics to guide image colorization. The rationale is that human beings perceive and distinguish colors based on the semantic categories of objects. Starting from an autoregressive model, we generate image color distributions, from which diverse colored results are sampled. We propose two ways to incorporate object semantics into the colorization model: through a pixelated semantic embedding and a pixelated semantic generator. Specifically, the proposed network includes two branches. One branch learns what the object is, while the other branch learns the object colors. The network jointly optimizes a color embedding loss, a semantic segmentation loss and a color generation loss, in an end-to-end fashion. Experiments on Pascal VOC2012 and COCO-stuff reveal that our network, when trained with semantic segmentation labels, produces more realistic and finer results compared to the colorization state-of-the-art.

Keywords Image colorization · Semantic segmentation · Pixelated semantics

1 Introduction

Color has been at the center stage of computer vision for decades (e.g., Swain and Ballard 1991; Comaniciu and Meer 1997; Pérez et al. 2002; Khan et al. 2009; van de Sande et al. 2010; Lou et al. 2015; Vondrick et al. 2018). Many vision challenges, including object detection and visual tracking, benefit from color (Khan et al. 2009, 2012; Danelljan et al. 2014; Vondrick et al. 2018). Consequently, color constancy (Gijssen et al. 2010) and color correction (Sanchez and Binefa 2000) methods may further enhance visual recog-

nition. Likewise, color is commonly added to gray-scale images to increase their visual appeal and perceptually enhance their visual content (e.g., Welsh et al. 2002; Iizuka et al. 2016; Zhang et al. 2016; Royer et al. 2017; Deshpande et al. 2017). This paper is about image colorization.

Human beings excel in assigning colors to gray-scale images since they can easily recognize the objects and have gained knowledge about their colors. No one doubts the sea is typically blue and a dog is never naturally green. Although many objects have diverse colors, which makes their prediction quite subjective, humans can get around this by simply applying a bit of creativity. However, it remains a significant challenge for machines to acquire both the world knowledge and “imagination” that humans possess.

Previous works in image colorization require reference images (Gupta et al. 2012; Liu et al. 2008; Charpiat et al. 2008) or color scribbles (Levin et al. 2004) to guide the colorization. Recently, several automatic approaches (Iizuka et al. 2016; Larsson et al. 2016; Zhang et al. 2016; Royer et al. 2017; Guadarrama et al. 2017) have been proposed based on deep convolutional neural networks. Despite the improved colorization, there are still common pitfalls that make the colored images appear less realistic. We show some examples in Fig. 1. The cases in (a) without semantics suffer from incorrect semantic understanding. For instance, the cow is

Communicated by Ling Shao, Hubert P. H. Shum, Timothy Hospedales.

✉ Jiaojiao Zhao
j.zhao3@uva.nl
Jungong Han
jungonghan77@gmail.com
Ling Shao
ling.shao@ieee.org
Cees G. M. Snoek
cgmsnoek@uva.nl

- ¹ Universiteit van Amsterdam, Amsterdam, The Netherlands
- ² University of Warwick, Coventry, UK
- ³ Inception Institute of Artificial Intelligence, Abu Dhabi, UAE

Fig. 1 Colorization without and with semantics generated using the network from this paper. We rescale all output images to their original proportions. **a** The method without semantics assigns unreasonable colors to objects, such as the colorful sky and the blue cow. The method

with semantics generates realistic colors for the sky (first column), the man (second column) and the cow (third column). **b** The method without semantics fails to capture long-range pixel interactions (Royer et al. 2017). With semantics, the model performs better

assigned a blue color. The cases in (b) without semantics suffer from color pollution. Our objective is to effectively address both problems to generate better colorized images with high quality.

Both traditional (Chia et al. 2011; Ironi et al. 2005) and recent colorization solutions (Larsson et al. 2016; Iizuka et al. 2016; He et al. 2016; Zhang et al. 2016, 2017) have highlighted the importance of semantics. However, they only explore image-level class semantics for colorization. As stated by Dai et al. (2016), image-level classification favors translation invariance. Obviously, colorization requires representations that are, to a certain extent, translation-variant. From this perspective, semantic segmentation (Long et al. 2015; Chen et al. 2018; Noh et al. 2015), which also requires translation-variant representations, provides more reasonable semantic guidance for colorization. It predicts a class label for each pixel. Similarly, according to Zhang et al. (2016) and Larsson et al. (2016), colorization assigns each pixel a color distribution. Both challenges can be viewed as an image-to-image prediction problem and formulated as a pixel-wise prediction task. We show several colorized examples after using pixelated semantic-guidance in Fig. 1a, b. Besides providing sharp boundaries which helps to prevent color bleeding, the color distributions of specific object types enforce additional constraints, which helps to alleviate the ambiguity in color recovery. Together, the fine-grained semantic information helps to precisely colorize specific objects.

In this paper, we study the relationship between colorization and semantic segmentation. Our proposed network is able to be harmoniously trained for semantic segmentation and colorization. By using such multi-task learning, we explore how pixelated semantics affects colorization. Differing from the preliminary conference version of this work (Zhao et al. 2018), we view colorization here as a sequential pixel-wise color distribution generation task, rather than a

pixel-wise classification task. We design two ways to exploit pixelated semantics for colorization, one by guiding a color embedding function and the other by guiding a color generator. Using these strategies, our methods produce diverse vibrant images on two datasets, Pascal VOC2012 (Everingham et al. 2015) and COCO-stuff (Caesar et al. 2018). We further study how colorization can help semantic segmentation and demonstrate that the two tasks benefit each other. We also propose a new quantitative evaluation method using semantic segmentation accuracy.

The rest of the paper is organized as follows. In Sect. 2, we introduce related work. Following, in Sect. 3, we describe the details of our colorization network using pixelated semantic guidance. Experiments and results are presented in Sect. 4. We conclude our work in Sect. 5.

2 Related Work

2.1 Colorization by Reference

Colorization using references was first proposed by Welsh et al. (2002), who transferred the colors by matching the statistic within the pixel's neighborhood. Rather than relying on independent pixels, Ironi et al. (2005) transferred colors from a segmented example image based on their observation that pixels with the same luminance value and similar neighborhood statistics may appear in different regions of the reference image, which may have different semantics and colors. Tai et al. (2005) and Chia et al. (2011) also performed local color transfer by segmentation. Bugeau et al. (2014) and Gupta et al. (2012) proposed to transfer colors at pixel level and super-pixel level. Generally, finding a good reference with similar semantics is key for this type of methods. Previously, Liu et al. (2008) and Chia et al. (2011) relied on image retrieval methods to choose good references. Recently, deep

learning has supplied more automatic methods in Cheng et al. (2015) and He et al. (2018). In our approach, we use a deep network to learn the semantics from data, rather than relying on a reference with similar semantics.

2.2 Colorization by Scribble

Another interactive way to colorize a gray-scale image is by placing scribbles. This was first proposed by Levin et al. (2004). The authors assumed that pixels nearby in space-time, which have similar gray levels, should have similar colors as well. Hence, they solved an optimization problem to propagate sparse scribble colors. To reduce color bleeding over object boundaries, Huang et al. (2005) adopted an adaptive edge detection to extract reliable edge information. Qu et al. (2006) colorized manga images by propagating scribble colors within the pattern-continuous regions. Yatziv and Sapiro (2006) developed a fast method to propagate scribble colors based on color blending. Luan et al. (2007), further extended (Levin et al. 2004) by grouping not only neighboring pixels with similar intensity but also remote pixels with similar texture. Several more current works (Zhang et al. 2017; Sangkloy et al. 2017) used deep neural networks with scribbles trained on a large dataset and achieved impressive colorization results. In all these methods, which use hints like strokes or points, provide an important means for segmenting an image into different color regions. We prefer to learn the segmentation rather than manually labelling it.

2.3 Colorization by Deep Learning

The earliest work applying a deep neural network was proposed by Cheng et al. (2015). They first grouped images from a reference database into different clusters and then learned deep neural networks for each cluster. Later, Iizuka et al. (2016) pre-trained a network on ImageNet for a classification task, which provided global semantic supervision. The authors leveraged a large-scale scene classification database to train a model, exploiting the class-labels of the dataset to learn the global priors. Both of these works treated colorization as a regression problem. In order to generate more saturated results, Larsson et al. (2016) and Zhang et al. (2016) modeled colorization as a classification problem. Zhang et al. (2016) applied cross-channel encoding as self-supervised feature learning with semantic interpretability. Larsson et al. (2016) claimed that interpreting the semantic composition of the scene and localizing objects were key to colorizing arbitrary images. Nevertheless, these works only explored image-level classification semantics. Our method takes the semantics one step further and utilizes finer pixelated semantics from segmentation.

Further, generative models have more recently been applied to produce diverse colorization results. Currently,

several works (Cao et al. 2017; Isola et al. 2017; Frans 2017) have applied a generative adversarial network (GAN) (Radford et al. 2016). They were able to produce sharp results but were not as good as the approach proposed by Zhang et al. (2016). Variational autoencoders (VAE) (Kingma and Welling 2014) have also been used to learn a color embedding (Deshpande et al. 2017). This method produced results with large-scale spatial co-ordination but tonelessness. Royer et al. (2017) and Guadarrama et al. (2017) applied PixelCNN (van den Oord et al. 2016; Salimans et al. 2017) to generate better results. We use PixelCNN as the backbone in this paper.

3 Methodology

In this section, we will detail how pixelated semantics improves colorization. We will first introduce our basic colorization backbone. Then, we will present two ways to exploit object semantics for colorization. Our network structure is summarized in Fig. 2.

3.1 Pixelated Colorization

To arrive at image colorization with pixelated semantics, we start from an autoregressive model. It colorizes each pixel conditioned on the input gray image and previously colored pixels. Specifically, a conditional PixelCNN (van den Oord et al. 2016) is utilized to generate per-pixel color distributions, from which we sample diverse colorization results.

We rely on the CIE **Lab** color space to perform the colorization, since it was designed to be perceptually uniform with respect to human color vision and only two channels **a** and **b** need to be learned. An image with a height H and a width W is defined as $X \in R^{H \times W \times 3}$. X contains n ($= H \times W$) pixels. In raster scan order: row by row and pixel by pixel within every row, the value of the i th pixel is denoted as X_i . The input gray-scale image, represented by light channel **L**, is defined as $X^L \in R^{H \times W \times 1}$. The objective of colorization is to predict the **a** and **b** channels $\hat{Y} \in R^{H \times W \times 2}$. Different from the **RGB** color space, **Lab** has the range $[0; 100] \times [-127; 128] \times [-127; 128]$.

To reduce computation and memory requirements, we prefer to produce color images with low resolution. This is reasonable since the human visual system resolves color less precisely than intensity (Van der Horst and Bouman 1969). As stated in Royer et al. (2017), image compression schemes, such as JPEG, or previously proposed techniques for automatic colorization also apply chromatic subsampling. The output images can be easily converted back to their original proportions. We can rescale the generated color channels and concatenate them with the original gray channel to produce the final colorized images with their original sizes.

Fig. 2 Pixelated semantic colorization. The three colored flows (arrows) represent three variations of our proposal. The purple flow illustrates the basic pixelated colorization backbone (Sect. 3.1). The purple flow combined with the blue flow obtains a better color embedding with more semantics (Sect. 3.2.1). The purple flow, blue flow and green

flow together define our final model, a pixelated colorization model conditioned on gray-scale image and semantic labels (Sect. 3.2.2). Here, f^θ is a color embedding function, h^ϕ is a semantic segmentation head and g^ω is the autoregressive generation model. There are three loss functions L_{seg} , L_{emb} and L_{gen} (Sect. 3.3)

By adopting PixelCNN for image colorization, a joint distribution with condition is modelled as van den Oord et al. (2016):

$$p(\hat{Y}|X^L) = \prod_{i=1}^n p(\hat{Y}_i|\hat{Y}_1, \dots, \hat{Y}_{i-1}, X^L). \quad (1)$$

All the elementary per-pixel conditional distributions are modelled using a shared convolutional neural network. As all variables in the factors are observed, training can be executed in parallel.

Furthermore, X^L can be replaced by a good embedding learned from a neural network. Taking g^ω as the generator function and f^θ as the embedding function, each distribution in Eq. (1) can be rewritten as:

$$p(\hat{Y}_i|\hat{Y}_1, \dots, \hat{Y}_{i-1}, X^L) = g_i^\omega(\hat{Y}_1, \dots, \hat{Y}_{i-1}, f^\theta(X^L)). \quad (2)$$

As the purple flow in Fig. 2 shows, there are two components included in our model. A deep convolutional neural network (f_θ) produces a good embedding of the input gray-scale image. Then an autoregressive model uses the embedding to generate a color distribution for each pixel. The final colorized results are sampled from the distributions using a pixel-level sequential procedure. We first sample $\hat{Y}_1$ from $p(\hat{Y}_1|X^L)$, then sample $\hat{Y}_i$ from $p(\hat{Y}_i|\hat{Y}_1, \dots, \hat{Y}_{i-1}, X^L)$ for all i in $\{2, \dots, n\}$.

3.2 Pixelated Semantic Colorization

Intuitively, semantics is the key to colorizing objects and scenes. We will discuss how to embed pixelated semantics in our colorization model for generating diverse colored images.

3.2.1 Pixelated Semantic Embedding

Considering the conditional pixelCNN model introduced above, a good embedding of the gray-scale image $f^\theta(X^L)$ greatly helps to generate the precise color distribution of each pixel. We first incorporate semantic segmentation to improve the color embedding. We use X^S to denote the corresponding segmentation map. Then, we learn an embedding of the gray-scale image conditioned on X^S . We replace $f^\theta(X^L)$ with $f^\theta(X^L|X^S)$. Thus, the new model learns the distribution in Eq. (2) as:

$$p(\hat{Y}_i|\hat{Y}_1, \dots, \hat{Y}_{i-1}, X^L, X^S) = g_i^\omega(\hat{Y}_1, \dots, \hat{Y}_{i-1}, f^\theta(X^L|X^S)). \quad (3)$$

Here the semantics only directly affects the color embedding generated from the gray-scale image, but not the autoregressive model.

Incorporating semantic segmentation can be straightforward, i.e., using segmentation masks to guide the colorization learning procedure. Such a way enables the training phase to directly obtain guidance from the segmentation masks, which clearly and correctly contain semantic information. However, it is not suitable for the test phase as segmentation masks are needed. Naturally, we can rely on an off-the-shelf segmentation model to gain segmentation masks for all the test images, but it is not elegant. Instead, we believe it is best to simultaneously learn the semantic segmentation and the colorization, making the two tasks benefit each other, as we originally proposed in Zhao et al. (2018).

Modern semantic segmentation can easily share low-level features with the color embedding function. We simply need to plant an additional segmentation branch h^ϕ following a few bottom layers, like the blue flow shown in Fig. 2. Specifically,

we adopt the semantic segmentation strategies from Chen et al. (2018). At the top layer, we apply atrous spatial pyramid pooling, which exploits multiple scale features by employing multiple parallel filters with different dilation rates. The final prediction ($h^\varphi(X^L)$) is the fusion of the features from the different scales, which helps to improve segmentation. The two tasks have different top layers for learning the high-level features. In this way, semantics is injected into the color embedding function. By doing so, a better color embedding with more semantic awareness is learned as input to the generator. This is illustrated in Fig. 2, by combining the purple flow and the blue flow.

3.2.2 Pixelated Semantic Generator

A good color embedding with semantics aids the generator to produce more correct color distributions. Furthermore, the generator is likely to be further improved with semantic labels. Here, we propose to learn a distribution conditioned on previously colorized pixels, a color embedding of gray-scale images with semantics ($f^\theta(X^L|X^S)$), and pixel-level semantic labels. We rewrite Eq. (3) as:

$$\begin{aligned} p(\hat{Y}_i | \hat{Y}_1, \dots, \hat{Y}_{i-1}, X^L, X^S) \\ = g_i^\omega(\hat{Y}_1, \dots, \hat{Y}_{i-1}, f^\theta(X^L|X^S), h^\varphi(X^L)). \end{aligned} \quad (4)$$

Intuitively, this method is capable of using semantics to produce more correct colors of objects and more continuous colors within one object. It is designed to address the two issues mentioned in Fig. 1. The whole idea is illustrated in Fig. 2 by combining the purple flow with the blue and green flows.

We consider two different ways to use pixelated semantic information to guide the generator. The first way is to simply concatenate the color embedding $f^\theta(X^L)$ and the segmentation prediction $h^\varphi(X^L)$ along the channels and then input the fusion to the generator. The second way is to apply a feature transformation introduced by Perez et al. (2018) and Wang et al. (2018). Specifically, we use convolutional layers to learn a pair of transformation parameters from the segmentation predictions. Then, a transformation is applied to the color embedding using these learned parameters. We find the first way works better. Results will be shown in the Experiment section.

3.3 Networks

In this section, we provide the details of the network structure and the optimization procedure.

Network Structure Following the scheme in Fig. 2, three components are included: the color embedding function f^θ , the semantic segmentation head h^φ and the autoregressive

model g^ω . Correspondingly, three loss functions are jointly learned, which will be introduced later. The three flows represent the three different methods introduced above. The purple flow illustrates the basic pixelated colorization. The purple flow combined with the blue flow results in the pixelated semantic embedding. The purple flow combined with the blue and green flows, results in the pixelated semantic generator.

Inspired by the success of the residual block (He et al. 2016; Chen et al. 2018) and following Royer et al. (2017), we apply gated residual blocks (van den Oord et al. 2016; Salimans et al. 2017), each of which has two convolutions with 3×3 kernels, a skip connection and a gating mechanism. We apply atrous (dilated) convolutions to several layers to increase the network's field-of-view without reducing its spatial resolution. Tables 1 and 2 list the details of the color embedding branch and the semantic segmentation branch, respectively. The gray rows are shared by the two branches.

Loss Functions During the training phase, we train the colorization and segmentation simultaneously. We try to minimize the negative log-likelihood of the probabilities:

$$\arg \min_{\theta, \varphi, \omega} \sum -\log p(\hat{Y} | f^\theta(X^L), h^\varphi(X^L)). \quad (5)$$

Specifically, we have three loss functions L_{emb} , L_{seg} and L_{gen} to train the color embedding, the semantic segmentation and the generator, respectively. The final loss function L_{sum} is the weighted sum of these loss functions:

$$L_{sum} = \lambda_1 * L_{emb} + \lambda_2 * L_{seg} + \lambda_3 * L_{gen}. \quad (6)$$

Following Salimans et al. (2017), we use discretized mixture logistic distributions to approximate the distribution in

Table 1 Color embedding branch structure

Module	Resolution	Channels	Dilation
Color embedding $f^\theta(X^L)$			
Convolution $3 \times 3/1$	128	64	–
Residual block $\times 2$	128	64	–
Convolution $3 \times 3/2$	64	128	–
Residual block $\times 2$	64	128	–
Convolution $3 \times 3/2$	32	256	–
Residual block $\times 2$	32	256	–
Convolution $3 \times 3/1$	32	512	–
Residual block $\times 3$	32	512	2
Convolution $3 \times 3/1$	32	512	–
Residual block $\times 3$	32	512	4
Convolution $3 \times 3/1$	32	160	–

Feature spatial resolution, number of channels and dilation rate are listed for each module. The italic rows indicate the bottom layers are shared with the semantic segmentation branch (detailed in Table 2)

Table 2 Semantic segmentation branch structure

Module	Resolution	Channels	Dilation
Semantic segmentation $h^\psi(X^L)$			
Convolution $3 \times 3/1$	128	64	–
Residual block $\times 2$	128	64	–
Convolution $3 \times 3/2$	64	128	–
Residual block $\times 2$	64	128	–
Convolution $3 \times 3/2$	32	256	–
Residual block $\times 2$	32	256	–
Convolution $3 \times 3/1$	32	512	–
Residual block $\times 3$	32	512	2
Convolution $3 \times 3/1$	32	512	–
Residual block $\times 3$	32	512	2
Convolution $3 \times 3/1$	32	#class	6
Convolution $3 \times 3/1$	32	#class	12
Convolution $3 \times 3/1$	32	#class	18
Add	32	#class	–

Feature spatial resolution, number of channels and dilation rate are listed for each module. #class means the number of semantic categories. The italic rows indicate the bottom layers are shared with the color embedding branch (detailed in Table 1)

Eqs. (3) and (4). A mixture of 10 logistic distributions is applied. Thus, both L^{emb} and L^{gen} are discretized mixture logistic losses.

As for semantic segmentation, generally it should be performed in the RGB image domain as colors are important for semantic understanding. However, the input of our network is a gray-scale image which is more difficult to segment. Fortunately, the network incorporating colorization learning supplies color information which in turn strengthens the semantic segmentation for gray-scale images. The mutual benefit among the three learning parts is the core of our network. It is also important to realize that semantic segmentation, as a supplementary means for colorization, is not required to be very precise. We use the cross entropy loss with the standard softmax function for semantic segmentation (Chen et al. 2018).

4 Experiments

4.1 Experimental Settings

Datasets We report our experiments on Pascal VOC2012 (Everingham et al. 2015) and COCO-stuff (Caesar et al. 2018). The former is a common semantic segmentation dataset with 20 object classes and one background class. Our experiments are performed on the 10,582 images for training and the 1449 images in the validation set for test-

Fig. 3 Color images, gray-scale images and segmentation maps from **a** Pascal VOC and **b** COCO-stuff. COCO-stuff has more semantic categories than Pascal VOC

ing. COCO-stuff is a subset of the COCO dataset (Lin et al. 2014) generated for scene parsing, containing 182 object classes and one background class on 9000 training images and 1000 test images. We train separate networks for Pascal VOC2012 and COCO-stuff. In order to reduce the computation and memory requirements, we rescale each input gray-scale image to 128×128 and produce the color maps with 32×32 , as shown in Table 1. The resolution of the color maps is 1/4 of the input image. Fig. 3 shows some examples

Fig. 4 Colorizations from the embedding functions f^θ using the purple flow and the purple-blue flow. **a** Colorization without semantic-guidance (first row) and with semantic-guidance (second row). With semantics, better colorizations are produced. **b** Visualization of the pre-

dicted **a** and **b** color channels of the colorizations. The top row shows the results without semantic-guidance and the bottom row with semantic-guidance. With semantics, the predicted colors have less noise and look more consistent

with natural scenes, objects and artificial objects from the datasets.

Implementation Commonly available pixel-level annotations intended for semantic segmentation are sufficient for our colorization method. We do not need new pixel-level annotations for colorization. We train our network with joint color embedding loss, semantic segmentation loss and generating loss with the weights $\lambda_1:\lambda_2:\lambda_3 = 1:100:1$, so that the three losses are similar in magnitude. Our multi-task learning for simultaneously optimizing colorization and semantic segmentation effectively avoids overfitting. The Adam optimizer (Kingma and Ba 2015) is adopted. We set an initial learning rate equal to 0.001, momentum to 0.95 and second momentum to 0.9995. We apply Polyak parameter averaging (Polyak and Juditsky 1992).

4.2 Effect of Segmentation on the Embedding Function f^θ

We first study how semantic segmentation helps to improve the color embedding function f^θ . Following the method introduced in Sect. 3.2.1, we jointly train the purple and blue flows shown in Fig. 2. In this case, the semantic segmentation branch only influences the color embedding function. To illustrate the effect of pixelated semantics, we compare the color embeddings generated from the embedding function f^θ in Fig. 4. Obviously, as can be seen, semantic-guidance enables better color embeddings. For example, the sky in the first picture looks more consistent, and the sheep are assigned reasonable colors. However, the results without semantic-guidance appear less consistent. For instance, there is color pollution on the dogs and the sheet in the second picture.

Fig. 5 Colorization from the generators g^o , when relying on the purple flow and the purple-blue flow. Examples from **a** Pascal VOC and **b** COCO-stuff are shown. For both datasets, the top row shows results

from the model without semantic-guidance and the bottom row shows the ones with semantic-guidance. The results with semantic-guidance have more reasonable colors and better object consistency

Further, in order to more clearly show the predicted color channels of the color embeddings, we remove the light channel L and only visualize the chrominances **a** and **b** in Fig. 4b. Interestingly, without semantic-guidance, the predicted colors are more noisy, as shown in the top row. However, with semantic-guidance, the colors are more consistent and echo the objects well. From these results, one clearly sees that colorization profits from semantic information. These comparisons support our idea and illustrate that pixelated semantics is able to enhance semantic understanding, leading to more consistent colorization.

In theory, we should obtain better colorization when a better color embedding is input into the generator. In Fig. 5, we show some final colorizations produced by the generator g^o . Our method using pixelated semantics works well on the two datasets. The results look more realistic. For instance, the fifth example in the Pascal VOC dataset is a very challenging case. The proposed method generates consistent and reasonable color for the earth even with an occluded object. For the last example in Pascal VOC, it is surprising that the horse bit is assigned a red color although it is very tiny. The proposed method processes details well. We also show vari-

Fig. 6 Incorporating pixelated semantics. Result comparisons between using **a** concatenation and **b** feature transformation to incorporate pixelated semantics. Concatenation generates more natural color images

ous examples from COCO-stuff, including animals, humans, fruits, and natural scenes. The model trained with semantics performs better. Humans are given normal skin color in the third and fifth examples. The fruits have uniform colors and look fresh.

Fig. 7 Colorizations generated by the embedding functions f^θ , using three variants of our network. The top row shows the results of the purple flow. The second row shows the results of the purple-blue flow.

The bottom row shows the results of the purple-blue-green flow. Each colorization is followed by the corresponding predicted chrominances. The purple-blue-green flow produces the best colorization

4.3 Effect of Segmentation on the Generator g^ω

In the next experiment, we add semantics to the generator as described in Sect. 3.2.2 (combining the purple flow with the blue and green flows). This means the generator produces a current pixel color distribution conditioned not only on the previous colorized pixels and the color embeddings from the gray image, but also on the semantic labels. We apply two different ways to incorporate semantics as introduced in Sect. 3.2.2. Using concatenation generates more natural colorful images than using feature transformation. Qualitative results are shown in Fig. 6. We prefer concatenation for the following experiments.

As we train the three loss functions L^{emb} , L^{seg} and L^{gen} simultaneously, we want to know whether the color embeddings produced by the embedding function are further improved. In Fig. 7, we compare the color embeddings generated by the embedding functions of the purple flow (shown in the top row), the purple-blue flow (shown in the second row) and the purple-blue-green flow (shown in the bottom row). Visualizations of color embeddings followed by the corresponding predicted chrominances are given. As can be seen, the addition of the green flow further improves the embedding function. From the predicted **a** and **b** visualizations, we observe better cohesion of colors for the objects. Clearly, the colorization benefits from the multi-task learning by jointly training the three different losses.

Indeed, using semantic labels as condition to train the generator results in better color embeddings. Moreover, the

final generated colorized results will also be better. In Fig. 8, we compare the results from the three methods: pixelated colorization without semantic guidance (the purple flow), pixelated semantic color embedding (the purple-blue flow), and pixelated semantic color embedding and generator (the purple-blue-green flow). The purple flow does not always understand the object semantic well, sometimes assigning unreasonable colors to objects, such as the cow in the third example of Pascal VOC, the hands in the second example and the apples in the last example of COCO-stuff. In addition, it also suffers from inconsistency and noise on objects. Using pixelated semantics to guide the color embedding function reduces the color noise and somewhat improves the results. Adding semantic labels to guide the generator improves the results further. As shown in Fig. 8, the purple-blue-green flow produces the most realistic and plausible results. Note that it is particularly apt at processing the details and tiny objects. For instance, the tongue of the dog is red and the lip and skin of the baby have very natural colors.

To conclude, these experiments demonstrate our strategies using pixelated semantics for colorization are effective.

4.4 Effect of Colorization on the Segmentation

From the previous discussion, it is concluded that semantic segmentation aids in training the color embedding function and the generator. The color embedding function and the generator also help each other. As stated in Sect. 3, the three

Fig. 8 Colorizations produced by the generators g^w , using three variants of our network on **a** Pascal VOC and **b** COCO-stuff: the purple flow (first row), the purple-blue flow (second row) and the purple-blue-green

flow (third row). Using pixel-level semantics to guide the generator in addition to the color embedding function achieves the most realistic results

learnings could benefit each other. Thus, we study whether colorization is able to improve semantic segmentation.

Color is Important for Semantic Segmentation As we observed in Zhao et al. (2018), color is quite critical for semantic segmentation since it captures some semantics. A

simple experiment is performed to stress this point. We apply the Deeplab-ResNet101 model (Chen et al. 2018) without conditional random field as post-processing, trained on the Pascal VOC2012 training set for semantic segmentation. We test three versions of the validation images, including gray-

Fig. 9 Segmentation results in terms of mean-IoU on gray-scale images, proposed colorized images and original color images, on the Pascal VOC2012 validation dataset. Color aids semantic segmentation

scale images, original color images and our colorized images. The mean intersection over union (mean-IoU) is adopted to evaluate the segmentation results. As seen in Fig. 9, with the original color information, the accuracy of 72.1% is much better than the 66.9% accuracy of the gray images. The accuracy obtained using our proposed colorized images is only 1.8% lower than using the original RGB images. This again demonstrates that our colorized images are realistic. More importantly, the proposed colorized images outperform the gray-scale images by 3.4%, which further supports the importance of color for semantic understanding.

Colorization Helps Semantic Segmentation In order to illustrate how colorization influences semantic segmentation, we train three semantic segmentation models on gray-scale images using our network structure: (1) we jointly train semantic segmentation and colorization; (2) we only train semantic segmentation from a pre-trained colorization model; (3) we only train semantic segmentation from scratch. We train all models on the training set of Pascal VOC 2012 and test them on the validation set. As validating loss reflects the semantic segmentation accuracy on the validation set, we compare the validating loss of the three models.

As seen in Fig. 10, the model trained on a pre-trained colorization model converges first. The loss is stable from the 18th epoch and the stable loss value is about 0.043. The model trained from scratch has the lowest starting loss but converges very slowly. Starting from the 55th epoch, the loss plateaus at 0.060. As expected, the pre-trained colorization model makes semantic segmentation achieve better accuracy. We believe the colorization model has already learned some semantic information from the colors, as also observed by Zhang et al. (2016). Further, our multi-task model jointly trained with semantic segmentation and colorization obtains the lowest validating loss of 0.030, around the 25th epoch. This supports our statement that the two tasks with the three loss functions are able to be learned harmoniously and benefit each other.

Fig. 10 Semantic segmentation validating loss comparisons. Three models are trained for 50 epochs. Training from a pre-trained colorization model is better than training from scratch. Jointly training obtains the lowest validating loss, which demonstrates colorization helps to improve semantic segmentation

Fig. 11 Samples diversity. Histogram of SSIM scores on the Pascal VOC validation dataset shows the diversity of the multiple colorized results. Some examples with their specific SSIM scores are also shown. Our model is able to produce appealing and diverse colorizations

4.5 Sample Diversity

As our model is capable of producing diverse colorization results for one gray-scale input, it is of interest to know whether or not pixelated semantics reduces the sample diversity. Following Guadarrama et al. (2017), we compare two outputs from the same gray-scale image with multiscale structural similarity (SSIM) (Wang et al. 2003). We draw the distribution of SSIM scores for all the compared pairs on the Pascal VOC validation dataset. As shown in Fig. 11, most of the output pairs have an SSIM score between 0.8 and 0.95. The examples shown in the figure demonstrate the pairs have the same content but different colors for details,

such as the eyes of the bird and the pants of the lady. Usually, the large backgrounds or objects with different colors in a pair of outputs cause lower SSIM scores. For instance, the backgrounds and birds in the first example. We believe pixelated semantics does not destroy the sample diversity. We will show more diverse colorization results in the next section.

4.6 Comparisons with State-of-the-Art

Generally, we want to produce visually compelling colorization results, which can fool a human observer, rather than recover the ground-truth. As discussed previously, colorization is a subjective challenge. Thus, both qualitative and quantitative evaluations are difficult. As for quantitative evaluation, some papers (Zhang et al. 2016; Iizuka et al. 2016) apply *Top-5* and/or *Top-1* classification accuracies after colorization to assess the performance of the methods. Other papers (He et al. 2018; Larsson et al. 2016) use the peak signal-to-noise ratio (PSNR), although it is not a suitable criteria for colorization, especially not for a method like ours, which produces multiple results. Also, color fidelity is suitable to be used for evaluating methods generating a single colorization result, which share a common goal to provide a color image closer to the original one. In Deshpande et al. (2015) and Larsson et al. (2016), the authors apply root mean squared error (RMSE) of the 2-channel images compared to the ground-truth to evaluate color fidelity. For qualitative evaluation, human observation is mostly used (Zhang et al. 2016; Iizuka et al. 2016; He et al. 2018; Royer et al. 2017; Cao et al. 2017).

In this paper, we propose a new evaluation method. We use semantic segmentation accuracy to assess the performance of each method, since we know semantics is key to colorization. This is more strict than classification accuracies. Specifically, we calculate the mean-IoU for semantic segmentation results from the colorized images. We use this procedure to compare our method with single colorization methods. For qualitative evaluation, we use the method from our previous work (Zhao et al. 2018). We ask 20 human observers, including research students and people without any image processing knowledge, to do a test on a combined dataset including the Pascal VOC2012 validation and the COCO-stuff subset. Given a colorized image or the real ground-truth image, the observers should decide whether it looks natural or not.

4.6.1 Single Colorization State-of-the-Art

We compare the proposed method with the single colorization state-of-the-art (Zhang et al. 2016; Iizuka et al. 2016; Larsson et al. 2016). In addition to the proposed semantic segmentation accuracy evaluation, we also report PSNR. We report RMSE of the 2 color channels **a** and **b** compared to the

Table 3 Quantitative evaluation

	Mean-IoU (%)	PSNR (dB)	RMSE
Iizuka et al. (2016)	67.6	24.20	10.66
Larsson et al. (2016)	68.8	24.56	10.26
Zhang et al. (2016)	68.1	22.81	12.37
Ours	70.3	23.15	11.43
Ground-truth (color)	72.1	NA	NA

Comparisons of semantic segmentation accuracies, and PSNRs and RMSEs between colorized results and the ground-truth, on the Pascal VOC validation set. Our method performs better according to the mean-IoU value

Bold values indicate the best performance in the comparisons

Italic values indicate the results achieved using ground-truth

Table 4 Qualitative evaluation

	Naturalness (%)
Single colorization	
Iizuka et al. (2016)	88.61
Larsson et al. (2016)	86.99
Zhang et al. (2016)	88.66
Diverse colorization	
Deshpande et al. (2017)	75.30
Cao et al. (2017)	80.00
Royer et al. (2017)	89.89
Ours	94.65
Ground-truth	99.58

Comparisons of the naturalness. Our colorizations are more natural than others

Bold values indicate the best performance in the comparisons

Italic values indicate the results achieved using ground-truth

ground truth. We use the Deeplab-ResNet101 model again for semantic segmentation. In this case, we only sample one result for each input, using our method.

Result comparisons are shown in Table 3. Our method has a lower PSNR than Iizuka et al. (2016) and Larsson et al. (2016). The comparisons of RMSE are similar with that of PSNR. Both depend on the ground-truth and over-penalize semantically plausible results with a colorization that differs from the ground-truth (He et al. 2018). Both Iizuka et al. (2016) and Larsson et al. (2016) obtain lower RMSEs because their objective is to minimize the distance between their outputs and ground-truth. However, our method outperforms all the others in semantic segmentation accuracy. This demonstrates that our colorizations are more realistic and contain more perceptual semantics.

For qualitative comparison, we report the naturalness of each method according to 20 human observations in Table 4. Three of the single colorization methods perform comparatively. Our results are more natural. Selected examples are shown in Fig. 12. The method by Iizuka et al. (2016) pro-

Fig. 12 Comparisons with single colorization state-of-the-art. Our results look more saturated and realistic

duces good results, but sometimes assigns unsuitable colors to objects, like the earth in the fourth example. The results from Larsson et al. (2016) look somewhat grayish. Zhang et al. (2016)'s method can generate saturated results but suffers from color pollution. Compared to these, our colorizations are spatially coherent and visually appealing. For instance, the color of the bird in the third example and the skin of the human in the last example, both look very natural.

4.6.2 Diverse Colorization State-of-the-Art

We also compare our method with the diverse colorization state-of-the-art (Royer et al. 2017; Cao et al. 2017; Deshpande et al. 2017). All of these are based on a generative model. We only qualitatively compare these by human observation. We use each model to produce three colorized samples. We report

the results in Table 4. Royer et al. (2017) apply PixelCNN to get natural images. Our results are even more natural. Several examples are shown in Fig. 13. Deshpande et al. (2017), using a VAE, generate sepia toned results. Cao et al. (2017), applying a GAN, output plausible results but with mixed colors. Royer et al. (2017) also produces saturated results but with color pollution. Our generated colored images have fine-grained and vibrant colors and look realistic.

4.6.3 Results on Legacy Black and White Photographs

We also colorize some legacy black and white photographs from renowned photographers Henri Cartier-Bresson and Ansel Adams, along with Thylacine which went extinct in 1936. Results are shown in Fig. 14. The model also works well on old black and white photos.

Fig. 13 Comparisons with diverse colorization state-of-the-art. The diverse results generated by our method look fairly good

4.7 Result Analysis

To further analyze the effect of object semantics on colorization, we quantify the colorization performance for each object class in the Pascal VOC validation set. We report

RMSEs per object in Fig. 15. Naturally, RMSE is highest on background as no specific object semantics is utilized for this class when we train our model. A person is the most difficult object to colorize as the colors of clothing are so diversiform. Low RMSEs for other objects, like bicycle and sheep, illus-

Fig. 14 Results on legacy black and white photos. The model also works well on old black and white photos

Fig. 15 Colorization results of per object class on the Pascal VOC validation set using RMSE. Without considering background, the model performs worse for persons as the clothing has too much diversity in color. The RMSEs are lower for all other objects. Our model helps to assign reasonable colors to objects

trate that incorporating semantics helps to precisely assign reasonable colors to objects.

4.8 Failure Cases

Our method is able to output realistic colorized images but it is not perfect. There are still some failure cases encountered by the proposed approach as well as other automatic systems. We provide a few failure cases in Fig. 16. Usually, it is highly challenging to colorize different kinds of food. They are artificial and variable. It is also difficult to learn the semantics of images containing several tiny and occluded objects. Moreover, our method cannot handle the

Fig. 16 Failure cases. Food, tiny objects and artificial objects are still very challenging

objects with unclear semantics. Although we exploit semantics for improving colorization, we do not have very many categories. We believe a finer semantic segmentation with more class labels will further enhance the results.

5 Conclusion

We propose pixelated semantic colorization to address a limitation of automatic colorization: object color inconsistency due to limited semantic understanding. We study how to effectively use pixelated semantics to achieve good colorization. Specifically, we design a pixelated semantic color embedding and a pixelated semantic generator. Both of these strengthen semantic understanding so that content confusion can be reduced. We train our network to jointly optimize colorization and semantic segmentation. The final colorized results on two datasets demonstrate the proposed strategies generate plausible, realistic and diverse colored images. Although we have achieved good results, our system is not perfect yet and has some challenges remaining. For instance, it cannot well process images with artificial objects, like food,

or tiny objects. More learning examples and finer semantic segmentation may further improve the colorization results in the future.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Bugeau, A., Ta, V. T., & Papadakis, N. (2014). Variational exemplar-based image colorization. *IEEE Transactions on Image Processing*, 23(1), 298–307.
- Caesar, H., Uijlings, J., & Ferrari, V. (2018). COCO-stuff: Thing and stuff classes in context. In *CVPR*.
- Cao, Y., Zhou, Z., Zhang, W., & Yu, Y. (2017). Unsupervised diverse colorization via generative adversarial networks. In *Joint European conference on machine learning and knowledge discovery in databases*.
- Charpiat, G., Hofmann, M., & Schölkopf, B. (2008). Automatic image colorization via multimodal predictions. In *ECCV*.
- Cheng, Z., Yang, Q., & Sheng, B. (2015). Deep colorization. In *ICCV*.
- Chen, L. C., Papandreou, G., Kokkinos, I., Murphy, K., & Yuille, A. L. (2018). Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4), 834–848.
- Chia, A. Y. S., Zhuo, S., Gupta, R. K., Tai, Y. W., Cho, S. Y., Tan, P., et al. (2011). Semantic colorization with internet images. *ACM Transactions on Graphics*, 30(6), 156.
- Comaniciu, D., & Meer, P. (1997). Robust analysis of feature spaces: Color image segmentation. In *CVPR*.
- Dai, J., Li, Y., He, K., & Sun, J. (2016). R-FCN: Object detection via region-based fully convolutional networks. In *NIPS*.
- Danelljan, M., Shahbaz Khan, F., Felsberg, M., & Van de Weijer, J. (2014). Adaptive color attributes for real-time visual tracking. In *CVPR*.
- Deshpande, A., Lu, J., Yeh, M. C., Chong, M. J., & Forsyth, D. (2017). Learning diverse image colorization. In *CVPR*.
- Deshpande, A., Rock, J., & Forsyth, D. (2015). Learning large-scale automatic image colorization. In *ICCV*.
- Everingham, M., Eslami, S. A., Van Gool, L., Williams, C. K., Winn, J., & Zisserman, A. (2015). The Pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1), 98–136.
- Frans, K. (2017). *Outline colorization through tandem adversarial networks*. [arXiv:1704.08834](https://arxiv.org/abs/1704.08834).
- Gijsenij, A., Gevers, T., & van de Weijer, J. (2010). Generalized gamut mapping using image derivative structures for color constancy. *International Journal of Computer Vision*, 86(2–3), 127–139.
- Guadarrama, S., Dahl, R., Bieber, D., Norouzi, M., Shlens, J., & Murphy, K. (2017). Pixcolor: Pixel recursive colorization. In *BMVC*.
- Gupta, R. K., Chia, A. Y. S., Rajan, D., Ng, E. S., & Zhiyong, H. (2012). Image colorization using similar images. In *Multimedia*.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *CVPR*.
- He, M., Chen, D., Liao, J., Sander, P. V., & Yuan, L. (2018). Deep exemplar-based colorization. *ACM Transactions on Graphics*, 37(4), 47.
- Huang, Y. C., Tung, Y. S., Chen, J. C., Wang, S. W., & Wu, J. L. (2005). An adaptive edge detection based colorization algorithm and its applications. In *Multimedia*.
- Iizuka, S., Simo-Serra, E., & Ishikawa, H. (2016). Let there be color!: Joint end-to-end learning of global and local image priors for automatic image colorization with simultaneous classification. *ACM Transactions on Graphics*, 35(4), 110.
- Ironi, R., Cohen-Or, D., & Lischinski, D. (2005). Colorization by example. In *Rendering techniques*.
- Isola, P., Zhu, J. Y., Zhou, T., & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In *CVPR*.
- Khan, F. S., Anwer, R. M., Van de Weijer, J., Bagdanov, A. D., Vanrell, M., & Lopez, A. M. (2012). Color attributes for object detection. In *CVPR*.
- Khan, F. S., Van De Weijer, J., & Vanrell, M. (2009). Top-down color attention for object recognition. In *ICCV*.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *ICLR*.
- Kingma, D. P., & Welling, M. (2014). Auto-encoding variational bayes. In *ICLR*.
- Larsson, G., Maire, M., & Shakhnarovich, G. (2016). Learning representations for automatic colorization. In *ECCV*.
- Levin, A., Lischinski, D., & Weiss, Y. (2004). Colorization using optimization. *ACM Transactions on Graphics*, 23(3), 689–694.
- Lin, T. Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., et al. (2014). Microsoft COCO: Common objects in context. In *ECCV*.
- Liu, X., Wan, L., Qu, Y., Wong, T. T., Lin, S., Leung, C. S., et al. (2008). Intrinsic colorization. *ACM Transactions on Graphics*, 27(5), 152.
- Long, J., Shelhamer, E., & Darrell, T. (2015). Fully convolutional networks for semantic segmentation. In *CVPR*.
- Lou, Z., Gevers, T., Hu, N., Lucassen, M. P., et al. (2015). Color constancy by deep learning. In *BMVC*.
- Luan, Q., Wen, F., Cohen-Or, D., Liang, L., Xu, Y. Q., & Shum, H. Y. (2007). Natural image colorization. In *Proceedings of the 18th eurographics conference on rendering techniques*.
- Noh, H., Hong, S., & Han, B. (2015). Learning deconvolution network for semantic segmentation. In *ICCV*.
- Pérez, P., Hue, C., Vermaak, J., & Gangnet, M. (2002). Color-based probabilistic tracking. In *ECCV*.
- Perez, E., Strub, F., De Vries, H., Dumoulin, V., & Courville, A. (2018). Film: Visual reasoning with a general conditioning layer. In *AAAI*.
- Polyak, B. T., & Juditsky, A. B. (1992). Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4), 838–855.
- Qu, Y., Wong, T. T., & Heng, P. A. (2006). Manga colorization. *ACM Transactions on Graphics*, 25(3), 1214–1220.
- Radford, A., Metz, L., & Chintala, S. (2016). Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*.
- Royer, A., Kolesnikov, A., & Lampert, C. H. (2017). Probabilistic image colorization. In *BMVC*.
- Salimans, T., Karpathy, A., Chen, X., Kingma, D. P., & Bulatov, Y. (2017). PixelCNN++: A PixelCNN implementation with discretized logistic mixture likelihood and other modifications. In *ICLR*.
- Sanchez, J. M., & Binefa, X. (2000). Improving visual recognition using color normalization in digital video applications. In *ICME*.
- Sangkloy, P., Lu, J., Fang, C., Yu, F., & Hays, J. (2017). Scribbler: Controlling deep image synthesis with sketch and color. In *CVPR*.
- Swain, M. J., & Ballard, D. H. (1991). Color indexing. *International Journal of Computer Vision*, 7(1), 11–32.
- Tai, Y. W., Jia, J. Y., & Tang, C. K. (2005). Local color transfer via probabilistic segmentation by expectation-maximization. In *CVPR*.
- van de Sande, K. E. A., Gevers, T., & Snoek, C. G. M. (2010). Evaluating color descriptors for object and scene recognition. *IEEE*

- Transactions on Pattern Analysis and Machine Intelligence*, 32(9), 1582–1596.
- Van der Horst, G. J., & Bouman, M. A. (1969). Spatiotemporal chromaticity discrimination. *Journal of the Optical Society of America*, 59(11), 1482–1488.
- van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., Graves, A., et al. (2016). Conditional image generation with PixelCNN decoders. In *NIPS*.
- Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., & Murphy, K. (2018). Tracking emerges by colorizing videos. In *ECCV*.
- Wang, Z., Simoncelli, E. P., & Bovik, A. C. (2003). Multiscale structural similarity for image quality assessment. In *The 37th Asilomar conference on signals, systems & computers*.
- Wang, X., Yu, K., Dong, C., & Loy, C. C. (2018). Recovering realistic texture in image super-resolution by deep spatial feature transform. In *CVPR*.
- Welsh, T., Ashikhmin, M., & Mueller, K. (2002). Transferring color to greyscale images. *ACM Transactions on Graphics*, 21(3), 277–280.
- Yatziv, L., & Sapiro, G. (2006). Fast image and video colorization using chrominance blending. *IEEE Transactions on Image Processing*, 15(5), 1120–1129.
- Zhang, R., Isola, P., & Efros, A. A. (2016). Colorful image colorization. In *ECCV*.
- Zhang, R., Zhu, J. Y., Isola, P., Geng, X., Lin, A. S., Yu, T., et al. (2017). Real-time user-guided image colorization with learned deep priors. In *SIGGRAPH*.
- Zhao, J., Liu, L., Snoek, C. G. M., Han, J., & Shao, L. (2018). Pixel-level semantics guided image colorization. In *BMVC*.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.