

DeepObjStyle: Deep Object-based Photo Style Transfer

Indra Deep Mastan and Shanmuganathan Raman
Indian Institute of Technology Gandhinagar
Gandhinagar, Gujarat, India
{indra.mastan, shanmuga}@iitgn.ac.in

Abstract

One of the major challenges of style transfer is the appropriate image features supervision between the output image and the input images (style and content). An efficient strategy would be to define an object map between the objects of the style and the content images. However, such a mapping is not well established when there are semantic objects of different types and numbers in the style and the content images. It also leads to content mismatch in the style transfer output, which could reduce the visual quality of the results. We propose an object-based style transfer approach, called DeepObjStyle, for the style supervision in the training data-independent framework. DeepObjStyle preserves the semantics of the objects and achieves better style transfer in the challenging scenario when the style and the content images have a mismatch of image features. We also perform style transfer of images containing a word cloud to demonstrate that DeepObjStyle enables an appropriate image features supervision. We validate the results using quantitative comparisons and user studies.

1. Introduction and Prior Work

Style transfer is an ill-posed problem that synthesizes a new image using the style and the content images [3]. The aim is to perform the image feature enhancement [9, 22, 21, 23, 24] of the content image. Style transfer output captures objects from the content image. The objects of the style transfer output get the style features from the style image. Style transfer when the content and the style images have the same set of semantic objects enables applications such as puppet control and domain translation [16]. Style transfer methods when the content and the style images have different sets of semantic objects are called arbitrary style transfer methods [5].

We classify deep convolutional neural network (CNN) based style transfer methods for simplicity as follows: (1) training data-based and (2) training data-independent methods. Training data-based style transfer methods use sam-

ples of the style and the content images to train the CNN. The training process would learn the style features from the style image and the content features from the content image [1]. Training-independent style transfer methods do not train a CNN with the samples of the style and the content images and mostly focus on designing the loss function for the deep features extracted using VGG19 [4, 14, 7, 5].

Training-data based methods that use adversarial loss require many samples of the style and the content images [6]. Zhang *et al.* have shown how to minimize the number of sample images by mixing the encoded latent representations of the style and the content images [26]. Mechrez *et al.* proposed contextual loss (CL), which trains CNN using the samples of the content images [16]. The stylization by Li *et al.* [11] uses Microsoft COCO dataset [12] to train the decoder network. WCT2 [25] uses images from Microsoft COCO dataset to improvise [11].

Training data-independent setup includes the style transfer of arbitrary images using the *gram* loss [4] proposed by Gatys *et al.*. Luan *et al.* improvise [4] for photo-realistic style transfer and geometric structure preservation [14]. They validated the quality of style transfer output by checking style spillover between dissimilar objects, also known as the *content mismatch* criterion. However, the content mismatch is inherent when the semantic objects in the style and the content images are of different types and numbers. Kolkin *et al.* proposed a new technique using optimal transport called STROTSS [7].

Image features transfer between contextually similar objects is a challenging task in training data-based and training data-independent setup [10]. The training data-based methods have shown a good qualitative performance, but they could be biased towards the images in training samples and might not generalize well to the new style and content images. The training data-independent setting is challenging because of the lack of object context information learned from the sample images.

Another challenge in style transfer is when the style and the content images have different types or different numbers of objects (Fig. 2). One strategy would be to merge




Figure 1: The figure shows style transfer in the challenging scenario of the mismatch of image features between style and content images. The content image is the face of a person and the style image is the scene of the buildings (top left corner). It could be observed that DeepObjStyle output images with better quality.

objects with similar classes (*e.g.*, grass and tree) to minimize the content mismatch [14]. However, the strategy above would not be useful when the semantic objects are of different classes (Sec. 2).

We address the style transfer challenges by an object-based style transfer approach and investigate how to utilize features from all the objects in the style and the content image while minimizing the content mismatch. We call our deep object-based photo-style transfer method as DeepObjStyle. The strategy is to distribute the image features considering the semantics of the image features in the style and the content images. The object-based approach aims to enable image features supervision based on the contextually similar regions of the output.

DeepObjStyle achieves good perceptual quality in the style transfer output in the presence of different numbers of objects present in the style and the content images and when there is a mismatch of image features between the style and the content images. We also investigate the approach by putting word cloud in the style and the content images. The word cloud would not be readable when the unrelated image features are merged with it. Therefore, it helps to investigate the preservation of the structure of the objects in the output.

DeepObjStyle supervises features by minimizing distortions and considering the object context of the objects of the style and the content images. We propose unmapped object loss and mapped object loss for the distribution of image features to the style transfer output.

Contributions. The major contributions of the work are as follows.

1. DeepObjStyle preserves the semantics of the content image in the style transfer output (Fig. 1).
2. We investigate the performance of DeepObjStyle in the following two challenging scenarios. Style transfer when there is a low similarity between the features of the style and the content images (Fig. 3 and Fig. 5). Style transfer where the style and the content images contain a word-cloud (Fig. 4).
3. DeepObjStyle provides feature transfer from all the objects of the style and the content images while preserving the object structure (Fig. 2, Fig. 4, Fig. 6, and Fig. 7).
4. We evaluated our results using reference-based quality assessment: Pieapp [18], and no-reference quality assessment: NIMA [8]. We also performed a user study to compare the style transfer outputs. We show that DeepObjStyle outperforms the relevant state-of-the-art methods (Sec. 4).

2. Background

Content Mismatch. We describe the content mismatch challenge of style transfer using an example as follows. Given a content image C and a reference style image S , the objective is to synthesize a new image O , which contains the content features from C and the style features from S . Also, suppose that the style image S has two semantic objects and the content image C has three semantic objects (Fig. 2-b). The style transfer output O gets the content features (*i.e.*, three objects) from C . The style features in these three objects in O are received from the two objects

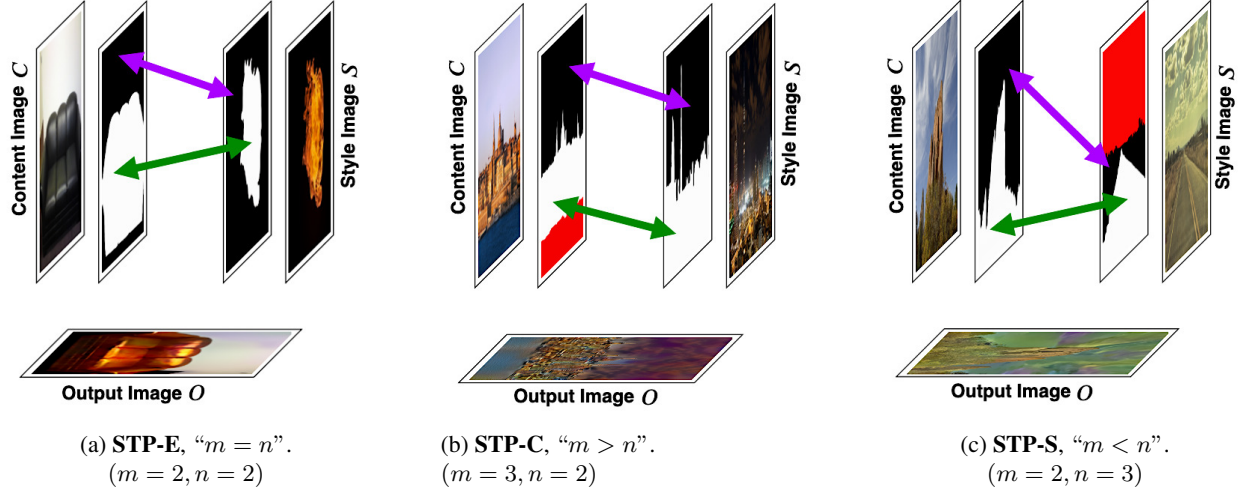


Figure 2: **Style Transfer Problems (STP)**. The figure illustrates STP. The style image S has n semantic objects, and the content image C has m semantic objects. The arrows show the object mapping and the red color in the segmentation mask shows the unmapped object. (a) STP-E: style transfer when the objects in S and C are equal. (b) STP-C: style transfer when C has more objects than S . (c) STP-S: style transfer when S has more objects than C . We describe STP in Sec. 2.

in S . The interesting thing to note is that providing style features from two objects to the three objects in O would introduce content mismatch as there would be one object in S , which transfer style features to multiple objects in O . The earlier style transfer methods do not address the challenging content mismatch scenario described above to the best of our knowledge. One simple strategy would be to discard style feature transfer to unmapped objects in O to minimize content mismatch.

Style Transfer Problems (STP). The task is to synthesize a new image O by using style features from the style image S and content features from the content image C . Suppose the content image C has m objects and the style image has n objects. We have shown Style Transfer Problems (STP) in Fig. 2. We discuss STP in detail as follows.

- **STP-E**: Style transfer when the objects in content image and the style image are equal (*i.e.*, $m = n$ in Fig. 2-a). Therefore, there exists a one-to-one object map between the objects of the style image and the objects of the content image. Moreover, STP-E is challenging when C and S have semantically different objects. It is due to the fact that the style transfer method has to synthesize a new image by mixing content features and style features of the different class of objects (*i.e.*, content mismatch in Fig. 5). We use a mapped object loss to transfer image features based on the object context for the above STP-E problem.
- **STP-C**: Style transfer when the content image has more semantic objects than that of the style image (*i.e.*, $m > n$ in Fig. 2-b). Therefore, a content mismatch occurs by the

utilization of style features from multiple objects of S to an object in O . We propose unmapped objects loss to minimize the content mismatch for the above style *diffusion* problem (style diffusion example in Fig. 6).

- **STP-S**: Style transfer when the style image has more objects than that of the content image (*i.e.*, $n > m$ in Fig. 2-c). Therefore, a content mismatch occurs due to the diffusion of style features from an object of S to multiple objects in O . We propose unmapped objects loss to minimize the effects of a content mismatch for the above style *utilization* problem (style utilization example in Fig. 7).

The style and the content features from C and S are extracted using the pre-trained VGG19 network [19] denoted by ϕ . We do not train CNN to learn image features from samples of the style and the content images. The extracted feature maps from ϕ are used to synthesize image features of the output O .

3. Our Approach

We have described that a good style transfer approach would allow features transfer while minimizing the effects of content mismatch and preserves the semantics of the objects in the style transfer output. One could perform style transfer without considering the objects, but it could result in low perceptual quality due to the mixing of features from different categories of objects in the output (Fig. 5).

Overview. We first define an one-to-one object map (OM) between the contextually similar objects of the style and the

content images using the segmentation mask. If the number of objects in the style and the content images are different, OM creates two categories of objects - mapped objects and unmapped objects. Mapped objects are contained in both the style and the content images and participate in OM. Unmapped objects are contained either in the style image or in the content image and do not belong to OM. Next, we use the mapped objects loss and the unmapped objects loss to distribute features from the input style image and content image to the style transfer output.

DeepObjStyle loss. We define DeepObjStyle loss \mathcal{L}_{dos} as a combination of mapped objects loss \mathcal{L}_{dos}^M and unmapped objects loss \mathcal{L}_{dos}^U in Eq. 1.

$$\mathcal{L}_{dos} = \alpha \mathcal{L}_{dos}^M + \beta \mathcal{L}_{dos}^U \quad (1)$$

Here, α and β are the coefficients. The mapped objects loss \mathcal{L}_{dos}^M uses segmentation masks to define an one-to-one object map between the objects of the style and the content images for style transfer. The unmapped objects loss \mathcal{L}_{dos}^U performs style transfer for unmapped objects independent of any mapping.

3.1. Mapped Objects Loss

The mapped objects loss \mathcal{L}_{dos}^M is computed between the output image, and the targeted style and the content images. \mathcal{L}_{dos}^M is a combination of the deep photo-style loss [14] denoted by \mathcal{L}_{dps} and contextual content loss denoted by $\mathcal{L}_{cl,C}$. We define \mathcal{L}_{dos}^M in Eq. 2.

$$\mathcal{L}_{dos}^M = \alpha_1 \mathcal{L}_{dps} + \alpha_2 \mathcal{L}_{cl,C} \quad (2)$$

Here, \mathcal{L}_{dps} transfers features from the style image S and the content image C to the output image O . Whereas, $\mathcal{L}_{cl,C}$ transfers the content features from C to O .

The deep photo-style loss \mathcal{L}_{dps} in Eq. 2 constrains image features of the style transfer output to be locally affine in colorspace to suppress distortions and yields photorealistic style transfer. The contextual content loss $\mathcal{L}_{cl,C}$ in Eq. 2 transfers the features by minimizing the dissimilarity between the contextually similar vectors sampled from the content representation $F_l[C]$ of C and the content representation $F_l[O]$ of O . Here, $F_l[\cdot]$ denotes the vectorized feature map present at layer l of features extractor $\phi(\cdot)$.

The contextual content loss $\mathcal{L}_{cl,C}$ at a layer l of ϕ is given as follows: $\mathcal{L}_{cl,C}^l = -\log CX(F_l[O], F_l[C])$. Here, $CX(\cdot, \cdot)$ computes the contextual similarity [16] between the content features of C and O . Thus, it helps to preserve the semantics of the objects in output O . We have described deep photo-style loss and contextual loss in the supplementary material in detail.

The mapped objects loss \mathcal{L}_{dos}^M defined in Eq. 2 is directly useful for STP-E as there are no unmapped objects (Fig. 2).

Similarly, the image features transfer for mapped objects in STP-S and STP-C is done using \mathcal{L}_{dos}^M .

3.2. Unmapped Objects Loss

Fig. 2 shows the unmapped objects in two scenarios: STP-C and STP-S. The unmapped objects loss \mathcal{L}_{dos}^U achieves style transfer for unmapped objects. \mathcal{L}_{dos}^U is computed between the output image O and the style image S .

\mathcal{L}_{dos}^U is a combination of the gram loss $\mathcal{L}_{gl,S}$, and contextual loss computed on style features, denoted by contextual style loss $\mathcal{L}_{cl,S}$ as shown in Eq. 3.

$$\mathcal{L}_{dos}^U = \beta_1 \sum_{l \in L} \mathcal{L}_{gl,S}^l + \beta_2 \sum_{l \in L} \mathcal{L}_{cl,S}^l \quad (3)$$

Here, $\mathcal{L}_{cl,S}^l$ works by minimizing the difference in the contextually similar vectors in the feature representations of the style image and the output image at the layer l . $\mathcal{L}_{gl,S}^l$ uses the gram matrix of the feature maps to spread features.

For STP-C, the contextual style loss is computed between the style image and the unmapped objects present in the output image to achieve *style diffusion*. For STP-S, the contextual style loss is computed between the unmapped objects in the style image S and the output image O to achieve *style utilization*.

Let $\vartheta_c = \{c_i\}_{i=1}^m$ be the set of segmentation channels in the content image C , where each segmentation channel represents an object of C . Similarly, let $\vartheta_s = \{s_i\}_{i=1}^n$ be the set of segmentation channels in the style image S . The style transfer for unmapped objects is described as follows.

Style diffusion (STP-C). It is style transfer when the number of semantic objects in the content image m is more than the number of semantic objects in the style image n (i.e., $m > n$ in Fig. 2-b). There are a total of $(m + n)$ objects in S and C , out of which $2n$ objects are mapped objects and $(m - n)$ unmapped objects. We describe the strategy for style diffusion for $(m - n)$ unmapped objects in the output O below.

First, we obtain the unmapped $(m - n)$ objects from O by computing the Hadamard product between O and the segmentation of unmapped $(m - n)$ objects $\vartheta_{c,n}$, where $\vartheta_{c,n} = \sum_{j=n}^m c_j$. Next, we provide the style features to the unmapped objects using the unmapped objects loss \mathcal{L}_{dos}^U introduced earlier in Eq. 3. For achieving this, the gram loss $\mathcal{L}_{gl,S}^l$ and the contextual style loss $\mathcal{L}_{cl,S}^l$ are defined in Eq. 4 and Eq. 5.

$$\mathcal{L}_{gl,S}^l = \frac{1}{2N_l^2} \sum_{ij} (G(F_l[O] \odot \vartheta_{c,n}) - G(F_l[S]))_{ij}^2 \quad (4)$$

Here, N_l denotes the filters in layer l . Gram matrix $G[\cdot]$ is an inner product between vectorized style feature maps taken from $F_l[\cdot]$, where $F_l[\cdot] \in \mathbb{R}^{N_l \times D_l}$ and D_l denotes the




Figure 3: This figure shows the style transfer in the presence of the mismatch of image features where the style image is not photo-realistic, but the content image is photo-realistic. It could be observed that DeepObjStyle suppress distortions and preserve the photo-realism of the content image in the style transfer output.

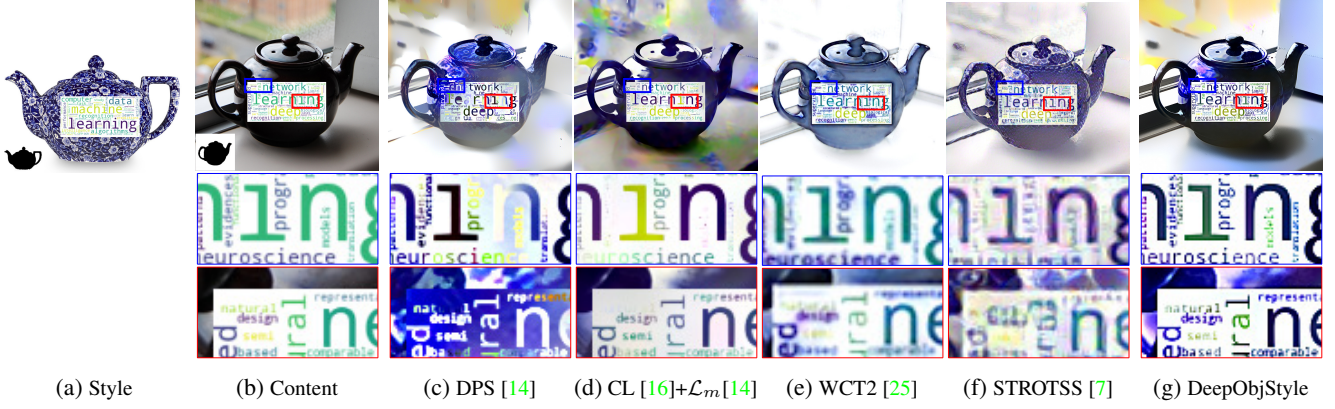


Figure 4: **Preservation of Structure (STP-E)**. The style image (a) and the content image (b) contains a word cloud. DPS [14] spills feature over the word cloud and suffers from the content mismatch. We integrated contextual loss [16] with photo-realism regularization \mathcal{L}_m [14]. It shows a better word cloud region, but the features for other areas are not well distributed. WCT2 [25] and STROTSS [7] do not preserve the structure and the word cloud is less readable. DeepObjStyle (ours) minimizes content mismatch and preserves the geometry of the objects (the images are best viewed after zooming).

size of the vectorized feature map related to the filter. Thus, $G[\cdot] = F_l[\cdot]F_l[\cdot]^T \in \mathbb{R}^{N_l \times N_l}$. Gram loss in Eq. 4 spreads style using the style representations $F_l[S]$ of the style image S and the style representations $F_l[O]$ of the output image O . Eq. 4 shows that we compute the feature correlations using the gram matrix of the style representation $G(F_l[S])$ and the representations of the unmapped objects of the output $G(F_l[O] \odot \vartheta_{c,n})$. The contextual style loss $\mathcal{L}_{cl,S}^l$ is given in Eq. 5.

$$\mathcal{L}_{cl,S}^l = -\log CX(F_l[O] \odot \vartheta_{c,n}, F_l[S]) \quad (5)$$

Here, $\mathcal{L}_{cl,S}^l$ works by minimizing the difference in the contextually similar vectors sampled from the style representation $F_l[S]$ of the style image S and the style representation $F_l[O] \odot \vartheta_{c,n}$ of the unmapped objects in the output image O .

Style utilization (STP-S). It is style transfer when the number of semantic objects in the style image is more than that

of the content image (*i.e.*, $n > m$ in Fig. 2-c). The challenge is to achieve style utilization of the $(n - m)$ unmapped objects in the style image using the loss \mathcal{L}_{dos}^U introduced earlier in Eq. 3.

First, we obtain the $(n - m)$ unmapped objects in the style image using their segmentation mask $\vartheta_{s,m}$, where $\vartheta_{s,m} = \sum_{j=m}^n s_j$. Next, we provide the style features to them using the gram loss $\mathcal{L}_{gl,S}^l$ and the contextual loss $\mathcal{L}_{cl,S}^l$ defined in Eq. 6 and Eq. 7.

$$\mathcal{L}_{gl,S}^l = \frac{1}{2N_l^2} \sum_{ij} (G(F_l[O]) - G(F_l[S] \odot \vartheta_{s,m}))_{ij}^2 \quad (6)$$

$$\mathcal{L}_{cl,S}^l = -\log CX(F_l[O], F_l[S] \odot \vartheta_{s,m}) \quad (7)$$

Eq. 6 and Eq. 7 show that we use the features from the unmapped objects in style utilization similar to style diffusion (Eq. 4 and Eq. 5). In the supplementary material, we provide more technical details for the unmapped objects.

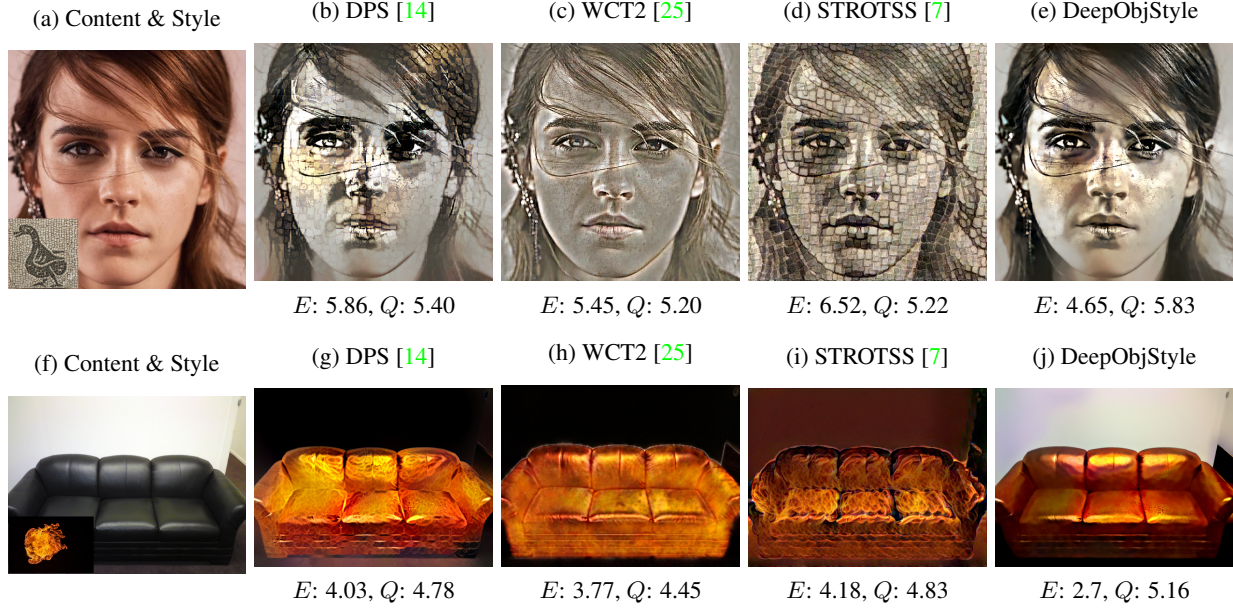


Figure 5: **Content Mismatch.** The style image is shown at the bottom left corner of the content image. These images have a content mismatch and challenging for style transfer [14]. The perceptual error score Pieapp [18] is denoted by E . The image quality score predicted by NIMA [8] is denoted by Q . DeepObjStyle achieves a minimum perceptual error score E and gets the highest quality score Q . DPS [14] and WCT2 [25] do not preserve the content features well. STROTSS [7] also suffers from the content mismatch. DeepObjStyle (ours) minimize the content mismatch and preserves the semantics of the objects.

4. Experimental Results

Fig. 1 and Fig. 3 show the photo-realistic style transfer when the style and the content images have different categories of the objects. We experiment with the content images that are photo-realistic and the task is to preserve the object structure properties. In Fig. 1, the style image is photo-realistic. In Fig. 3, the style image is not photo-realistic. The main intuition is that images with less deformation are more visually appealing. Neural style [4] deforms the geometry of the image. DPS [14] preserves object boundaries, but the style features are not distributed well. WCT2 [25] does not preserve fine image feature details in the style transfer output. STROTSS [7] does not distribute image features well. DeepObjStyle is able to preserve the photo-realistic properties of the content image in both the scenarios shown in Fig. 1 and Fig. 3.

Fig. 4 shows the style transfer when the style and content images contain a word cloud. The challenge here is to supervise style and content features while maintaining the readability of the text. DPS [14] spills-over unrelated features on the word-cloud. To investigate the photo-realistic style supervision with contextual loss CL [16], we integrate the photo-realism regularization module \mathcal{L}_m [14] with CL [16]. The photo-realism mostly suppresses distortions and preserves the structure of the objects in the output. \mathcal{L}_m [14]

with CL [16] does not distribute image features well. It might be because contextually similar features between the source and the target images were not well used in the output. WCT2 [25] and STROTSS [7] does not preserve the local level image features details and reduce the text readability. DeepObjStyle provides a better distribution of features even when the segmentation mask does not provide the position of the word cloud¹.

Fig. 5 shows the style transfer when the style and the content images have an extreme mismatch of image features. DPS [14], WCT2 [25], and STROTSS [7] output images with higher perceptual error and lower image quality score. DPS [14] does not preserve the semantics of the objects. WCT2 [25] trains a decoder on sample images. Therefore, the higher perceptual error might be because of the bias toward the sample images and lacking the generalization to the new images. STROTSS [7] transport style features onto the content image with minimum distortion to the geometry of the objects, but in the challenging scenario of content mismatch, the structure preservation reduces. DeepObjStyle outperforms other methods and preserves the semantics of the objects in the output.

Fig. 6 shows style transfer when the content image has more objects than that of the style image. It is worth noting that merging lake object with sky or building would result

¹The extended version of Fig. 4 is shown in the supplementary material.



Figure 6: **Style Diffusion (STP-C)**. The content image has more semantic objects than that of style image. The semantic objects of style and content images are shown by the segmentation mask at the bottom left corner. The sky and building are the mapped objects. The unmapped content object is the lake shown by red color in segmentation mask in (b). The geometry of objects is not clear in Neural style [4]. DPS [14] does not provide the style features to the unmapped lake object, thus style diffusion not achieved. WCT2 [25] and STROTSS [7] do not distribute image features well for the sky. DeepObjStyle (ours) achieves style diffusion while preserving the geometry of the objects (images are best viewed after zooming).




Figure 7: **Style Utilization (STP-S)**. The style image S has more semantic objects than that of the content image C . We map the grass in S with the sky in C to show the effects of content mismatch. DPS [14] does not achieve style utilization as the style features of the sky are not present in the output image. We integrated unmapped object loss in DPS and called it DPS^+ , for a fair comparison. DPS^+ achieves style utilization as we could see the color of the sky from S , but object boundaries are not visible. The image features are not much clear for WCT2 [25] and STROTSS [7]. DeepObjStyle achieves style utilization as the color of the sky is also taken from the sky of S . DeepObjStyle learns the object context and preserves the geometry of the objects in the output (see the cropped images).

in a high content mismatch. Neural style [4] spreads the style features disregarding the object boundaries. DPS [14] does not supervise features of the sky and the lake object well. WCT2 [25] does not fully utilize style features well, but it preserved better image features than STROTSS [7]. DeepObjStyle spreads the style features to all the objects in the output image while preserving the structure.

Fig. 7 shows the influence of content mismatch in style transfer when style image has more objects than that of content image. We purposefully mapped the sky of content image with trees of style image to make a challenging scenario for style transfer. DPS [14] does not preserve the content features well. To get a fair comparison, we improve DPS with the unmapped objects loss in DPS^+ method (we describe DPS^+ in the supplementary material). It utilizes features from all the semantic objects of the style and the content images. However, the preservation of the object structure is not achieved. WCT2 [25] and STROTSS [7] performs better but local image features details are not preserved. It can be observed that DeepObjStyle achieves style utilization while preserving the structure of the object.

5. Quantitative Comparison

Fig. 8 shows the no-reference quality assessment using NIMA [8] for 100 style transfer instances. NIMA [8] predicts the image quality score. The average image quality scores are as follows. Neural Style [4]: 5.13, DPS [14]: 5.23, WCT2 [25]: 5.35, STROTSS [7]: 4.88, and DeepObjStyle: 5.49. DeepObjStyle outperforms other methods by max image quality score.

Fig. 9 shows the reference-based quality assessment using perceptual error scores Pieapp [18] to investigate the distortion of the content features. The average Pieapp [18] scores are as follows. Neural Style [4]: 4.26, DPS [14]: 3.92, WCT2 [25]: 3.22, STROTSS [7]: 4.21, and DeepObjStyle: 2.83. DeepObjStyle outperforms other methods by minimum perceptual error.

We have also conducted a user study to validate the style transfer results. We took 18 style transfer instances and displayed them in random order. Each subject is asked to vote for the better-looking image. We collected feedback from 60 human experts with a total of 1080 votes. The number of votes for the methods is as follows. Neural Style [4]: 72,




Figure 8: **Image Quality Comparison.** The figure shows the comparison of image quality of style transfer output using NIMA [8]. DeepObjStyle achieves a better quality score as compared to other methods.




Figure 9: **Perceptual Error Comparison.** The figure shows perceptual error comparison with the reference of input content image using Pieapp [18]. DeepObjStyle shows minimum distortion in the output as compared to other methods. Therefore, output images of DeepObjStyle are more visually appealing.


DPS [14]: 78, WCT2 [25]: 174, STROTSS [7]: 76, and DeepObjStyle: 680 (highest votes).

The user study, perceptual error score [18], and image quality score [8] confirms the observation that DeepObjStyle transfer image features with better structure preservation and lesser distortions are visually more appealing².

6. Limitations

Fig. 10 and Fig. 11 show the limitations of DeepObjStyle. Fig. 10 shows the style transfer when style and con-

²We used the implementation of Neural style provided in [20], Tensorflow implementation of DPS given in [13], contextual loss implementation in [15], STROTSS implementation in [17], and WCT2 implementation in [2]. We have provided more visual comparisons and implementation details of our method in the supplementary material.



(a) Content & Style (b) STROTSS [7] (c) DeepObjStyle

Figure 10: **Limitation-1.** DeepObjStyle preserves a better object context, but the distribution of features might be improved further.



(a) Content & Style (b) STROTSS [7] (c) DeepObjStyle

Figure 11: **Limitation-2.** DeepObjStyle minimizes the distortion. Therefore, it misses fine texture information. STROTSS [7] does not preserve photo-realism. A similar effect is shown in the second row of Fig. 5.

tent images have an extreme mismatch of features. Fig. 11 shows the style transfer when the style image is an artistic image with painting style deformations. The limitations are due to the DeepObjStyle strategy of contextually similar feature comparison and minimizing deformations. We observed that mixing uncorrelated image features in the style transfer output results in a high perceptual error and low-quality output in general (Sec. 5).

7. Conclusion

DeepObjStyle achieves good image features supervision in many challenging scenarios, such as content mismatch and style transfer of images containing a word cloud. DeepObjStyle achieves style diffusion and style utilization while minimizing content mismatch. DeepObjStyle preserves the semantics of the objects and the object structure. The extensive experiments and quality assessment shows that DeepObjStyle outperforms relevant style transfer methods. We believe that perceptual quality in the extreme content mismatch scenario could be enhanced further. We propose designing a method that performs style transfer in the extreme content mismatch scenario as future work.

Acknowledgements. Indra Deep Mastan was supported by Visvesvaraya Ph.D. fellowship. Shanmuganathan Raman was supported by SERB Core Research Grant and SERB IMPRINT 2.

References

- [1] Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, and Gang Hua. Stylebank: An explicit representation for neural image style transfer. In *Proc. CVPR*, volume 1, page 4, 2017. 1
- [2] Clovaai. <https://github.com/clovaai/WCT2>, 2019. 8
- [3] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576*, 2015. 1
- [4] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016. 1, 2, 5, 6, 7
- [5] Shuyang Gu, Congliang Chen, Jing Liao, and Lu Yuan. Arbitrary style transfer with deep feature reshuffle. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8222–8231, 2018. 1
- [6] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2017. 1
- [7] Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. Style transfer by relaxed optimal transport and self-similarity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10051–10060, 2019. 1, 2, 5, 6, 7, 8
- [8] Christopher Lennan, Hao Nguyen, and Dat Tran. <https://github.com/ideal0/image-quality-assessment>, 2018. 2, 6, 7, 8
- [9] Chongyi Li, Jichang Guo, Fatih Porikli, and Yanwei Pang. Lightnet: a convolutional neural network for weakly illuminated image enhancement. *Pattern Recognition Letters*, 104:15–22, 2018. 1
- [10] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. In *Advances in Neural Information Processing Systems*, pages 386–396, 2017. 1
- [11] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. A closed-form solution to photorealistic image stylization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 453–468, 2018. 1
- [12] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 1
- [13] Yang Liu. <https://github.com/LouieYang/deep-photo-styletransfer-tf>, 2017. 8
- [14] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6997–7005, 2017. 1, 2, 4, 5, 6, 7, 8
- [15] Roey Mechrez. <https://github.com/roimehrez/contextualLoss>, 2018. 8
- [16] Roey Mechrez, Itamar Talmi, and Lihi Zelnik-Manor. The contextual loss for image transformation with non-aligned data. *European Conference on Computer Vision (ECCV)*, 2018. 1, 4, 5, 6
- [17] Nkolkin. <https://github.com/nkolkin13/STROTSS>, 2019. 8
- [18] Ekta Prashnani, Hong Cai, Yasamin Mostofi, and Pradeep Sen. Pieapp: Perceptual image-error assessment through pairwise preference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1808–1817, 2018. 2, 6, 7, 8
- [19] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [20] Cameron Smith. <https://github.com/cysmith/neural-style-tf>, 2016. 8
- [21] Jing-Wein Wang, Ngoc Tuyen Le, Jiann-Shu Lee, and Chou-Chen Wang. Color face image enhancement using adaptive singular value decomposition in fourier domain for face recognition. *Pattern Recognition*, 57:31–49, 2016. 1
- [22] Tianren Wang, Teng Zhang, and Brian C Lovell. Ebit: Weakly-supervised image translation with edge and boundary enhancement. *Pattern Recognition Letters*, 138:534–539, 2020. 1
- [23] Junyi Xie, Hao Bian, Yuanhang Wu, Yu Zhao, Linmin Shan, and Shijie Hao. Semantically-guided low-light image enhancement. *Pattern Recognition Letters*, 138:308–314, 2020. 1
- [24] Shibai Yin, Yibin Wang, and Yee-Hong Yang. A novel image-dehazing network with a parallel attention block. *Pattern Recognition*, 102:107255, 2020. 1
- [25] Jaesun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. Photorealistic style transfer via wavelet transforms. In *International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 5, 6, 7, 8
- [26] Yexun Zhang, Ya Zhang, and Wenbin Cai. Separating style and content for generalized style transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, 2018. 1