



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/190004/>

Version: Accepted Version

Proceedings Paper:

Ioannou, E. and Maddock, S. (2022) Depth-aware neural style transfer using instance normalization. In: Turner, M. and Vangorp, P., (eds.) Computer Graphics & Visual Computing (CGVC) 2022. Computer Graphics & Visual Computing (CGVC) 2022, 15-16 Sep 2022, Virtual conference. Eurographics Digital Library.

© 2022 The Author(s). Eurographics Proceedings © 2022 The Eurographics Association. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Depth-aware Neural Style Transfer using Instance Normalization

E. Ioannou¹ and S. Maddock¹

Department of Computer Science, The University of Sheffield, UK



Figure 1: Results using our approach: The colour and texture patterns of the style image are captured while depth information is preserved.

Abstract

Neural Style Transfer (NST) is concerned with the artistic stylization of visual media. It can be described as the process of transferring the style of an artistic image onto an ordinary photograph. Recently, a number of studies have considered the enhancement of the depth-preserving capabilities of the NST algorithms to address the undesired effects that occur when the input content images include numerous objects at various depths. Our approach uses a deep residual convolutional network with instance normalization layers that utilizes an advanced depth prediction network to integrate depth preservation as an additional loss function to content and style. We demonstrate results that are effective in retaining the depth and global structure of content images. Three different evaluation processes show that our system is capable of preserving the structure of the stylized results while exhibiting style-capture capabilities and aesthetic qualities comparable or superior to state-of-the-art methods. Project page: <https://ioannoue.github.io/depth-aware-nst-using-in.html>.

CCS Concepts

• **Computing methodologies** → **Image processing**; **Image representations**; • **Applied computing** → **Fine arts**; **Media arts**;

1. Introduction

Neural Style Transfer (NST) is concerned with the artistic stylization of various forms of data, such as images, videos and 3D models. In the context of 2D image stylization, which is where NST has predominantly been applied, this can be described as the process of transferring the style of one image onto an input ‘content’ image. The technique, which has attracted wide attention in academia and industry, is capable of mapping the style patterns of an artistic image onto an ordinary photograph, synthesizing a novel image that preserves the contents of the photograph while embodying the artistic influences of the particular artwork.

Many studies have extended the seminal 2D work of Gatys *et al.* [GEB16] to other media, such as 3D images, videos, and games [RDB17; HWL*17; GGZY18; DGV20]. Multiple studies have also addressed the computational complexity, speed or aesthetics and the visual quality of the stylized results [JAF16; ULVL16;

SKLO18; HJL*20]. Recently, a considerable amount of research has been devoted to enhancing the structure and depth-preserving capabilities of the NST algorithms based on the observation that the stylized images often neglect much of the content information by applying the style patterns evenly throughout the whole image [LCLR17; CLW*19; KKM19]. To eliminate these undesired effects, which are especially visible when the input content images include objects at various depths, algorithms have been proposed that in addition to content and style loss also encompass a depth reconstruction loss in training [LCLR17; CLW*19]. This is achieved by utilizing state-of-the-art depth estimation approaches [CFYD16].

We present an approach that is based on the image transformation network introduced by Johnson *et al.* [JAF16] and produces stylized results that preserve the global structure and depth of the contents. Our algorithm uses Instance Normalization (IN) layers in-

stead of Batch Normalization (BN), a modification to style transfer approaches proposed by Ulyanov *et al.* [UVL16] that improves the quality of the results. In addition, we utilize a state-of-the-art depth estimation network [RLH*20; RBK21] for the computation of the depth information that serves as an additional loss function to the content and style losses. This method extends the method by Liu *et al.* [LCLR17] which initially introduced the idea of incorporating depth reconstruction loss for the training of the image transformation network as a way of generating stylized results that take into account the depth of the input content images. We show that our approach, by making use of a more accurate depth estimation approach than Liu *et al.*, and also by replacing the Batch Normalization layers with Instance Normalization, is capable of producing results that better embody the style contrast across different areas of the image, thus improving upon the quality of the stylization.

The rest of the paper is organized as follows. Section 2 presents the related work. Our method is introduced and analyzed in detail in Section 3. Section 4 contains the results of our approach along with a discussion about its effectiveness. Finally, Section 5, supplies conclusions and discussion about future work.

2. Related work

Empowered by the comparable-to-human capabilities of Convolutional Neural Networks (CNNs) in object recognition, Gatys *et al.* suggested a system that reproduces famous paintings on natural images [GEB16]. The algorithm takes as input a content image and a style image and initializes a noise image which is subsequently optimized by seeking to minimize an objective function that encompasses definitions of content loss and style loss. The content is represented by the higher-level features of a pre-trained VGG-19 network [SZ15] while the style is considered as a set of summary statistics. For the representation of the style, features are extracted from multiple layers and feature correlations are calculated with the use of Gram matrices. Despite the effectiveness and sophistication of this procedure, it requires a notable amount of time to generate a single stylized image. The method of Johnson *et al.* [JAF16] avoids this slow optimisation process and proposes an algorithm that utilizes perceptual loss functions to train their networks. A generative model is optimized offline allowing the stylized output to be produced with a single forward pass, which is orders of magnitude faster. Similar algorithms manage to improve upon the speed and overall computational cost by learning feed-forward networks [ULVL16; UVL17], while further work considers the incorporation of multiple styles per model [DSK16; CYL*17] or even an arbitrary style per model [HB17; CS16; GCLY18].

Further improvements to earlier approaches were shown by using instance normalization (IN) – or contrast normalization – instead of batch normalization [UVL17]. Similarly to Johnson *et al.* [JAF16], Ulyanov *et al.* [ULVL16] use a generator that is composed of convolutions, pooling, upsampling and batch normalization. Ulyanov *et al.* [UVL16] suggest the same configuration but with contrast normalization layers in order to prevent the stylized results from depending on the contrast of the content image. The results demonstrate better quality and inhibit the undesired distribution of style patterns across the whole image [UVL17].

Other approaches view the problem from a different perspective,

trying to redefine what is considered to be the “style” of an artwork. Such attempts include the algorithm of Sanakoyeu *et al.* [SKLO18] which trains the network to focus on the details that are relevant for the style when measuring the similarity in content between the input and the stylized image, resulting in a more generalized procedure that avoids fixed style representations (captured by the features of a pre-trained VGG network). Another method presented by Hu *et al.* [HJL*20] aims to control the aesthetic of the stylized result by separately manipulating the colour and texture features. Their system inputs two different images that reference colour and texture instead of only one style image. These approaches seem to exceed the limitation of using a single style image and improve upon the aesthetic qualities of the stylized results, offering a viewpoint that, unlike previous studies, focuses on the style image’s characteristics and overall aesthetics.

Most of the aforementioned studies neglect to consider depth preservation and coherence of details, yet, depth information is considered to be of great importance when evaluating the visual quality of the NST methods’ results [JYF*19]. To address this limitation, Liu *et al.* [LCLR17] proposed a system that is based on the work of Johnson *et al.* [JAF16] but integrates depth estimation. Their method suggests the addition of depth reconstruction loss in training of the transformation network and makes use of a single-image depth perception network [CFYD16]. An extension to this work was implemented by Cheng *et al.* [CLW*19], whose approach focuses on retaining or enhancing the structure of the artistically stylized result. Using a global structure extraction network (represented by the depth map) and a local structure refinement network (represented by the image edges), they provide an adjustable way to control the amount of structure that is preserved when stylizing an image. This results in stylized outputs that do not suffer from style textures being scattered over the whole image and disrupting the content structures, which is aesthetically better, especially when the content image contains a face or multiple objects at various depths. However, as the authors suggest, the method might be unsuitable for users that prefer a more abstract feeling. Another approach by Kitov *et al.* [KKM19], which is based on the method of Huang *et al.* [HB17], applies stylization in different regions of the content image with different strength depending on their distance to the camera. Our approach focuses on retaining the global structure of an image, building on the initial method of Liu *et al.* [LCLR17].

Previous approaches rely on depth estimation methods to calculate depth from a 2D image. Predicting depth from a single RGB image is a long-standing problem in computer vision. The earliest data-driven methods that emerged with the rise of Deep Learning mostly make use of neural networks that are trained on ground-truth metric depth.

The recent work by Ranftl *et al.* [RLH*20; RBK21], relies on the idea that effective monocular depth estimation is tightly dependent on the variety and diversity of the training data. The method consists of a supervised model that is trained on five different diverse datasets, taking into account indoor and outdoor scenes including static and dynamic objects in various contexts. As part of our contributions, we compare these state-of-the-art methods on depth es-

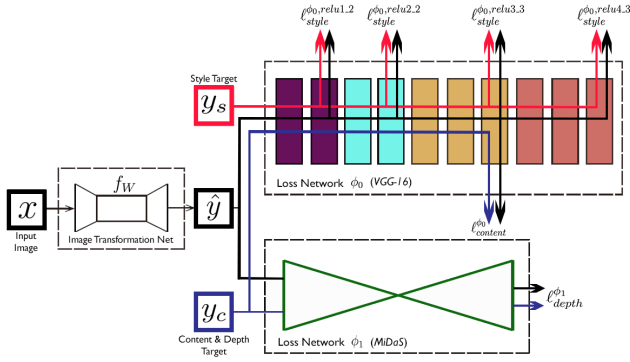


Figure 2: System Overview: The Image Transformation Network which transforms an input image x into an output image \hat{y} is trained. A pre-trained image classification network (VGG-16) is used to define the content and style losses. For the definition of depth loss, we use a state-of-the-art depth estimation network (MiDaS). To generate a style transfer model, for a particular style, the total loss (content loss + style loss + depth loss) is being optimized during training.

timation and utilize a depth estimation network for the training of our models that incorporates a depth reconstruction loss.

3. Method

3.1. Overview

Figure 2 provides an overview of the overall architecture used in our approach. Similarly to Liu *et al.* [LCLR17], our method uses an image transformation network (f_W) – a deep residual convolutional network – that transforms an input image x into an output image \hat{y} via the mapping $\hat{y} = f_W(x)$. Unlike Liu *et al.* [LCLR17], we use Instance Normalization (IN) layers instead of Batch normalization, so that normalization is applied to single images instead of a whole batch of images. This is based on the observation of Ulyanov *et al.* [UVL16] that such a modification makes the network agnostic to the contrast of the original images by preventing instance-specific mean and covariance shift. The final configuration of our network thus consists of (i) the model used by Johnson *et al.* [JAF16], but with IN layers, and (ii) a depth perception network that is used to capture the depth loss. Figure 3 demonstrates the improvements in the aesthetic of the stylized results generated when the network proposed by Johnson *et al.* is configured with IN layers instead of Batch Normalization. The original method of Johnson *et al.* discards most of the content information, applying style patterns evenly throughout the whole 2D image, whereas replacing Batch Normalization with IN favours structure and content preservation. Therefore, it is sensible to configure our architecture with IN layers since we aim to pay attention to depth and structure information and produce results with better style contrast.

3.2. Content & Style Losses

In addition to the image transformation network, we also use two loss networks to capture three different losses: a pre-trained image




Figure 3: Comparison of the results of the system proposed by Johnson *et al.* [JAF16] when using Batch Normalization (BN) and Instance Normalization (IN)

classification network to capture content loss and style loss, and a depth prediction network to capture depth loss. As with Johnson *et al.*, we use VGG-16 [SZ15] and its high-level features in order to define the content and style losses. Based on the observation that the deeper layers of a pre-trained convolutional network transform the input image into feature maps that increasingly care about the content of the image rather than any detail about the texture or colour of pixels, the content loss is defined by the squared Euclidean distance between the feature representations of the content image and the transformed image at a particular layer of the network ($relu2_2$):

$$l_{content}^{\phi_0}(\hat{y}, x) = \frac{1}{C_j H_j W_j} \|\phi_0^j(\hat{y}) - \phi_0^j(x)\|_2^2 \quad (1)$$

where ϕ_0 is the image classification network and ϕ_0^j represents the activations of the j^{th} layer of ϕ_0 when processing an image with shape $H \times W \times C$ where H denotes the height, W the width and C the number of channels.

For the calculation of the style loss, features are extracted from multiple layers and the feature correlations are given by the Gram matrix G which contains non-localized information about the image:

$$G_j^{\phi_0}(x)_{C,C'} = \frac{1}{C_j H_j W_j} \sum_{h=1}^{H_j} \sum_{w=1}^{W_j} \phi_0^j(x)_{h,w,c} \phi_0^j(x)_{h,w,c'} \quad (2)$$

The style loss is defined by the squared Frobenius norm between the Gram-based style representations of the transformed image \hat{y} and style image y :

$$l_{style}^{\phi_0,j}(\hat{y}, y) = \|G_j^{\phi_0}(\hat{y}) - G_j^{\phi_0}(y)\|_F^2 \quad (3)$$

The total style loss is then defined as:

$$l_{style}^{\phi_0}(\hat{y}, y) = \sum_{j \in J} l_{style}^{\phi_0, j}(\hat{y}, y) \quad (4)$$

where $J = \{relu1_1, relu2_2, relu3_3, relu4_3\}$ is the set of selected layers.

3.3. Depth Loss

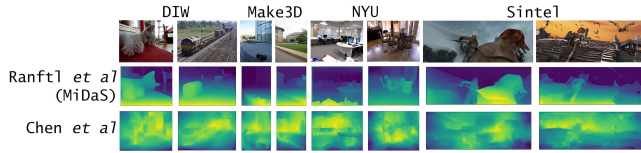


Figure 4: Visual qualitative comparison of the depth estimation methods by Ranftl et al. (MiDaS) [RLH*20; RBK21] and Chen et al. [CFYD16]

The superiority of *MiDaS* [RLH*20; RBK21] over the pre-existing single-image depth estimation methods guided our choice for the depth network that is utilized in order to compute the depth loss. In their work, they report an overall better performance in comparison with the system proposed by Chen *et al.* [CFYD16]. Figure 4 provides a visual comparison between the two methods applied on a variety of images. The images are sampled from a set of diverse datasets, including *DIW* [CFYD16], *NYU* [SHKF12], *Make3D* [SSN08] and *Sintel* [BWSB12]. We use images from the test datasets which were not used during the training of either of the two algorithms.

The chosen depth estimation network (ϕ_1) takes as input an image and directly calculates the depth map. The depth loss is thus defined as the Euclidean distance between the responses of the depth estimation network in regard to the original content image and the transformed image:

$$l_{depth}^{\phi_1}(\hat{y}, x) = \frac{1}{C_j H_j W_j} \|\phi_1(\hat{y}) - \phi_1(x)\|_2^2 \quad (5)$$

3.4. Training Details

The algorithm is trained on the Microsoft COCO dataset [LMB*14] which consists of 80k images. For the training, each image is resized to 256×256 . We use Adam optimizer [KB14] with a learning rate of 1×10^{-3} and train with a batch size of 4. As discussed previously, the content loss is computed at the *relu2_2* layer and the style reconstruction loss at layers *relu1_2*, *relu2_2*, *relu3_2* and *relu4_3* of the *VGG-16* loss network. The depth reconstruction loss is computed at the output layer of the *MiDaS* [RLH*20; RBK21] network. We found the optimal weights for the content, style and depth loss to be 1×10^5 , 1×10^{10} and 1×10^3 , respectively. The source code is available from the project's webpage: <https://ioannou.github.io/depth-aware-nst-using-in.html>.

4. Results and Discussion

We compare the results of our method against state-of-the-art approaches [GEB16; JAF16; LCLR17] both qualitatively and quantitatively. We perform a side-to-side visual comparison, we present the results of a user study designed to capture the subjectivity when evaluating the aesthetics of the results, and also consider quantitative metrics that assess the capability of the algorithms to preserve the content's depth and enhance the style contrast across the image.

4.1. Comparison with state-of-the-art methods




Figure 5: Visual comparison between the methods by Gatys et al. [GEB16], Johnson et al. [JAF16] (with IN), Liu et al. [LCLR17] and ours. The results of the methods by Gatys et al. [GEB16] and Johnson et al. [JAF16] were reproduced based on the original implementations of the authors whereas the results of Liu et al. [LCLR17] are retrieved directly from their paper with the authors' permission.

Figure 5 presents a visual side-by-side comparison between the results of our method and the results of the methods by Gatys *et al.* [GEB16], Johnson *et al.* [JAF16] (with instance normalization (IN)) and Liu *et al.* [LCLR17]. The images generated using the system by Gatys *et al.* [GEB16] depict the style patterns of the style images quite well, nevertheless, they fail to preserve the overall structure and 3D layout of the content image. The algorithm of Johnson *et al.* [JAF16] enhanced with Instance Normalization

does better in terms of preserving the contents of the image and not applying the style patterns evenly throughout the whole image. The same applies for the method of Liu *et al.* [LCLR17]. Although the results are quite similar, our system manages to better distinguish the objects that are located further away from the objects that are located near the camera. This is more visible in the last two images where our algorithm is capable of identifying the objects at the centre of the image and stylizing them appropriately, avoiding the distribution of uneven brush strokes in the background.




Figure 6: Illustration of our results for different content images and different painting styles. Our approach captures the colour and texture patterns of the style image, while retaining depth information, allowing the objects located at the centre of the image to stand out.

More results from models trained on various style images using our approach are demonstrated in Figure 6. While capturing the colour and texture patterns of the style image, our method also manages to retain the depth and allow the object located at the centre of the image to stand out. This is prominent in the first row of the figure and the third where the church and dandelion head, respectively, are being emphasized.

As Liu *et al.* [LCLR17] have noticed, there is a trade-off between perceptual loss and depth loss. If style loss is emphasized, the style patterns can be captured more accurately, whereas increasing the depth loss weight will preserve more of the contents and spatial layout of the image (Figure 7).




Figure 7: Increasing the depth weight results in more structure being preserved but it captures less of the style patterns. The style, content and the stylized result with no depth information are displayed on the left. Our results are shown in increasing depth weight from top left to the bottom. The top left image is generated with the lowest weight and the bottom right with the highest weight for the depth loss.

4.2. User study

To quantitatively gauge the aesthetic effect of our approach, we conducted a user study. We selected 5 different style images, including the common artistic paintings that are presented amongst previous NST studies' results, and 6 different content images that vary in form, colour, and content, ranging from landscapes to nature photographs and face portraits. The 6 content images are displayed in the left column of Figure 6, with the 5 style images and resulting stylized content images shown in the middle and right columns (with one of the style images repeated). The participants were shown a series of 30 sets of images. One of the images was generated by our algorithm whereas the other three were generated using previous algorithms (Gatys *et al.* [GEB16], Johnson *et al.* [JAF16] and Liu *et al.* [LCLR17]). The order that the images were shown was randomized. For the first half of the questions (1-15), the content and style images were not shown at the start of the question. The same set of 15 questions (16-30) was also presented (in random order) but with the content and style images that were used to generate the results revealed. We reasoned that excluding the content and style images of the questions would allow us to better evaluate the aesthetic effect of the final results, regardless of the generation process. The participants were asked to select the one image of the four stylized images that they visually preferred (i.e. their favourite stylization). We collected results from 20 participants. Examples of the questions shown to the participants can be found on the project's website: <https://ioannou.github.io/depth-aware-nst-using-in.html#userStudy>.

The results of the user study are shown in Table 1. The table shows the ratio of times each method was chosen as the most preferable method or amongst the most preferable methods, and also includes the overall vote distribution for all questions. Our method dominates the user preferences when the content and style images

Most preferable method ratio				
Content & Style	[GEB16]	[JAF16]	[LCLR17]	ours
Omitted	13.33%	13.33%	6.67%	80%
Revealed	60%	0%	33.33%	13.33%
Total votes				
Content & Style	[GEB16]	[JAF16]	[LCLR17]	ours
Omitted	22%	16.33%	18.67%	43%
Revealed	40.33%	15.33%	22.33%	22%

Table 1: The results of the user study. Stylized images using our method are compared against the methods of Gatys et al. [GEB16], Johnson et al. [JAF16] (with IN) and Liu et al. [LCLR17]. User preferences for the two different types of questions: omitting the content and style images from the question (omitted) and revealing this information (revealed). The first part of the table (Most preferable method ratio) shows the ratio of times a method is chosen as the most preferable method. There are occasions where more than one method is selected as the favourite of all participants – two or more methods collected the highest (and same) amount of votes. The second part of the table (Total votes) shows the ratio of the overall amount of votes each method has collected for all the questions

are not revealed to the participants whereas the method of Gatys et al. [GEB16] is the most popular method when the content and style images are included as part of the question. Although the seminal work of Gatys et al. produces stylizations that capture the style patterns more effectively, our method still performs well and it is superior to the other methods when comparing the final results. This also highlights the significance of the evaluation followed by each NST study, and how each aspect is valued. In this instance, we demonstrate that evaluating only the aesthetics of the results without considering the exact NST operation, the participants prefer different methods compared to when the content and style images are revealed. This is better illustrated by the graph in Figure 8. This shows our method is preferred for the majority of Questions 1-15 where the content and style images are not revealed, perhaps suggesting that when the style image is not shown, depth becomes an important factor in considering an image’s quality.

4.3. Metrics

Our third evaluation process makes use of metrics inspired by Liu and Zhu’s [LZ21] evaluation procedure. The aim here is to consider depth and global structure preservation. In terms of structural evaluation of an image, Liu and Zhu suggest Structural Similarity (SSIM) is more compatible with the human visual system (HVS) than peak signal-to-noise ratio (PSNR) and mean-square error (MSE), and thus more suitable for assessing the similarity between a content image and a stylized result. We also use histogram (Hist), average Hash (aHash), and difference hash (dHash) to compare our results with other state-of-the-art methods. The histogram can detect the tonal and colour intensity differences, whereas the image hash [Buc21] algorithms analyse the image structure on luminance and can be exploited to identify similar inputs. We use the method of Cai et al. [BX18] to perform decolourization on the im-

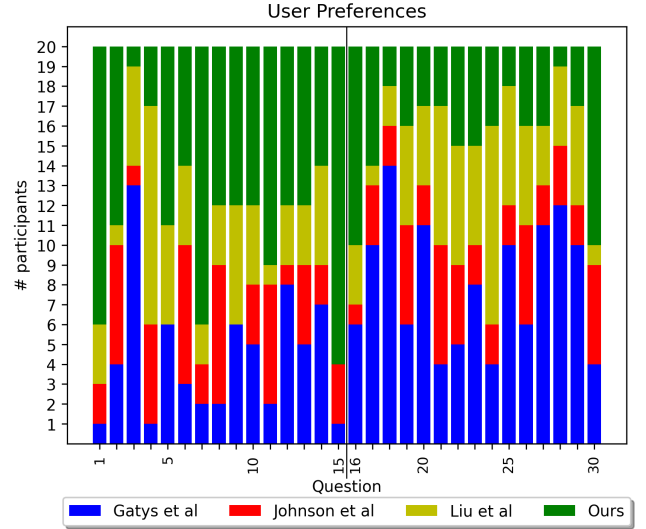


Figure 8: Overall user preferences for all the 30 questions. For questions 1-15 content and style images are omitted whereas for questions 16-30 the content and style images are shown as part of the question.

ages before computing the results. This process helps remove any colour (transforming RGB images to grayscale) while preserving the content information. This analysis is demonstrated in Table 2. Our method performs better in preserving the structure of the image and the method of Gatys et al. [GEB16] does better in preserving the tonal intensity differences, since their procedure is initialized with a copy of the original content image (which is iteratively being optimized).

	[GEB16]	[JAF16]	[LCLR17]	ours
SSIM	0.3951	0.4857	0.4176	0.4905
Hist	0.5518	0.5115	0.4555	0.4516
aHash	0.9219	0.8264	0.8316	0.8403
dHash	0.8767	0.7587	0.7378	0.7830

Table 2: The average values for SSIM, Hist, aHash and dHash metrics for the methods by Gatys et al. [GEB16], Johnson et al. [JAF16] (with IN), Liu et al. [LCLR17] and our method. The values are measured by comparing the stylized result of each method against the original content image. This is performed for 9 different pairs of images for each method and the averages are calculated. The best results are highlighted with red and the second best with cyan.

Additionally, we perform depth map and saliency map comparison. Ideally, the stylized result preserves more of the depth information and structure of the content image. We use the method of Ranftl et al. [RLH*20; RBK21] to compute the depth maps. We measure the structural similarity (SSIM) between the original image’s depth map and the stylized result’s depth map. Our method manages to preserve more of the depth information.

	[GEB16]	[JAF16]	[LCLR17]	ours
Depth map	0.8673	0.8848	0.8801	0.9112
Saliency map	0.4605	0.4723	0.4675	0.5021

Table 3: The average SSIM between the depth maps and saliency maps of the original image and the methods by Gatys et al. [GEB16], Johnson et al. [JAF16] (with IN), Liu et al. [LCLR17] and ours. The values are measured by comparing the depth and saliency map of the stylized result of each method against the depth and saliency map of the original content image. This is performed for 10 different pairs of images for each method and the averages are calculated. The best results are highlighted with red and the second best with cyan.

Saliency detection is considered an instance of image segmentation and can be used to identify visually predominant regions. The aim of our stylization method is to induce as minor change as possible to the saliency map of the content image, again resulting in more detail preservation. We repeat the structural similarity measurements on the saliency maps to give a more accurate estimate of the results. The method by Jiang et al. [JWY*13] is used to perform saliency detection. The average results for the depth and saliency maps structural similarity are reported in Table 3. Our method performs the best (highlighted with red) in preserving the details of the content as the depth and saliency map of our stylized results are closer to the original image’s.

Using an advanced depth estimation method, we show that the results can be significantly improved in comparison with the previous method of Liu et al. [LCLR17]. The global structure, depth, and dominant regions are being preserved resulting in aesthetically enhanced results.

For the depth map comparison, we used the method of Ranftl et al. [RLH*20; RBK21], whereas in the works of Liu et al. [LCLR17] and Cheng et al. [CLW*19] the method of Chen et al. [CFYD16] was used. Similarly to these methods, the same depth prediction network utilised during training was also used to drive the depth map comparison of the generated results. Future work could consider comparisons using both depth estimation algorithms ([RLH*20; CFYD16]).

4.4. Discussion

Visual side-by-side comparisons indicate that, similarly to the state-of-the-art methods, our system can properly capture colour and texture patterns of the style image, and in addition, it can produce stylizations that retain depth and allow the main object at the centre of the image stand out. The user study suggests that the effect our algorithm achieves has some positive impact on the aesthetics of the results as it is favoured by the participants. It also raises questions regarding the design of the user study and the form of its presentation, i.e. does showing the style image as part of the question lead a user’s thinking about what to consider when interpreting a picture’s aesthetic quality?

Lastly, we have provided a quantitative evaluation based on

particular metrics, capable of assessing the depth and structure-preserving capabilities of a method (depicted in Tables 2 and 3). A small set of images was used for this. A larger evaluation dataset will be considered in future work to increase the validity of the results. Additionally, although SSIM has been chosen as the metric to drive the comparison in depth and saliency maps, different metrics, such as MSE and PSNR, could also be considered.

These three evaluation approaches – visual side-by-side comparisons, user study and quantitative evaluation – are all commonly used in previous research, yet the best approach for evaluation is still an open question.

5. Conclusions

We have developed an approach for depth-aware neural style transfer on images. We have demonstrated that our models can effectively stylize 2D images while retaining the depth and global structure of the input content image. Our system replaces the Batch Normalization layers of the generator network with Instance Normalization and employs an advance depth prediction network for the calculation of a depth reconstruction loss. We have shown that Instance Normalization, which is not used in previous depth-aware style transfer methods, improves the quality of the results. In addition, we have shown that a more accurate depth estimation network can help maintain better style contrast across the image and further improve upon the preservation of the hierarchy and depth information.

We have evaluated our results both qualitatively and quantitatively using three different approaches, each of which has been used in previous studies. No single approach prevails. Whilst attempts have been made to propose robust quantitative evaluation procedures that do not rely on user studies [YTBF18], it remains an open question how best to evaluate stylized images. In our future work, we intend to consider the field of computational aesthetics assessment [ZMY21] and what this might offer to the evaluation process for NST techniques.

6. Acknowledgements

This research was supported by the EPSRC [grant number EP/R513313/1].

References

- [Buc21] BUCHNER, JOHANNES. *ImageHash*. 2021. URL: <https://pypi.org/project/ImageHash/> 6.
- [BWSB12] BUTLER, D. J., WULFF, J., STANLEY, G. B., and BLACK, M. J. “A naturalistic open source movie for optical flow evaluation”. *European Conf. on Computer Vision (ECCV)*. Ed. by A. FITZGIBBON ET AL. (EDS.) Part IV, LNCS 7577. Springer-Verlag, Oct. 2012, 611–625 4.
- [BX18] BOLUN CAI, XIANGMIN XU and XING, XIAOFEN. “Perception Preserving Decolorization”. *IEEE International Conference on Image Processing (2018)* 6.
- [CFYD16] CHEN, WEIFENG, FU, ZHAO, YANG, DAWEI, and DENG, JIA. “Single-image depth perception in the wild”. *arXiv preprint arXiv:1604.03901* (2016) 1, 2, 4, 7.

- [CLW*19] CHENG, MING-MING, LIU, XIAO-CHANG, WANG, JIE, et al. "Structure-preserving neural style transfer". *IEEE Transactions on Image Processing* 29 (2019), 909–920 1, 2, 7.
- [CS16] CHEN, TIAN QI and SCHMIDT, MARK. "Fast patch-based style transfer of arbitrary style". *arXiv preprint arXiv:1612.04337* (2016) 2.
- [CYL*17] CHEN, DONGDONG, YUAN, LU, LIAO, JING, et al. "Stylebank: An explicit representation for neural image style transfer". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, 1897–1906 2.
- [DGV20] DELIOT, THOMAS, GUINIER, FLORENT, and VANHOEY, KENNETH. *Real-time style transfer in Unity using deep neural networks*. 2020. URL: <https://blogs.unity3d.com/2020/11/25/real-time-style-transfer-in-unity-using-deep-neural-networks/> 1.
- [DSK16] DUMOULIN, VINCENT, SHLENS, JONATHON, and KUDLUR, MANJUNATH. "A learned representation for artistic style". *arXiv preprint arXiv:1610.07629* (2016) 2.
- [GCLY18] GU, SHUYANG, CHEN, CONGLIANG, LIAO, JING, and YUAN, LU. "Arbitrary style transfer with deep feature reshuffle". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018, 8222–8231 2.
- [GEB16] GATYS, LEON A, ECKER, ALEXANDER S, and BETHGE, MATTHIAS. "Image style transfer using convolutional neural networks". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, 2414–2423 1, 2, 4–7.
- [GGZY18] GAO, CHANG, GU, DERUN, ZHANG, FANGJUN, and YU, YIZHOU. *ReCoNet: Real-time Coherent Video Style Transfer Network*. 2018. arXiv: [1807.01197](https://arxiv.org/abs/1807.01197) [cs.CV] 1.
- [HB17] HUANG, XUN and BELONGIE, SERGE. "Arbitrary style transfer in real-time with adaptive instance normalization". *Proceedings of the IEEE International Conference on Computer Vision*. 2017, 1501–1510 2.
- [HJL*20] HU, ZHIYUAN, JIA, JIA, LIU, BEI, et al. "Aesthetic-Aware Image Style Transfer". *Proceedings of the 28th ACM International Conference on Multimedia*. 2020, 3320–3329 1, 2.
- [HWL*17] HUANG, HAOZHI, WANG, HAO, LUO, WENHAN, et al. "Real-time neural style transfer for videos". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, 783–791 1.
- [JAF16] JOHNSON, JUSTIN, ALAHI, ALEXANDRE, and FEI-FEI, LI. "Perceptual losses for real-time style transfer and super-resolution". *European conference on computer vision*. Springer. 2016, 694–711 1–7.
- [JWY*13] JIANG, HUAIZU, WANG, JINGDONG, YUAN, ZEJIAN, et al. "Salient Object Detection: A Discriminative Regional Feature Integration Approach". *2013 IEEE Conference on Computer Vision and Pattern Recognition*. 2013, 2083–2090. DOI: [10.1109/CVPR.2013.271](https://doi.org/10.1109/CVPR.2013.271) 7.
- [JYF*19] JING, YONGCHENG, YANG, YEZHOU, FENG, ZUNLEI, et al. "Neural style transfer: A review". *IEEE transactions on visualization and computer graphics* (2019) 2.
- [KB14] KINGMA, DIEDERIK P and BA, JIMMY. "Adam: A method for stochastic optimization". *arXiv preprint arXiv:1412.6980* (2014) 4.
- [KKM19] KITOV, VICTOR, KOZLOVTSSEV, KONSTANTIN, and MISHUSTINA, MARGARITA. "Depth-Aware Arbitrary Style Transfer Using Instance Normalization". *arXiv preprint arXiv:1906.01123* (2019) 1, 2.
- [LCLR17] LIU, XIAO-CHANG, CHENG, MING-MING, LAI, YU-KUN, and ROSIN, PAUL L. "Depth-aware neural style transfer". *Proceedings of the Symposium on Non-Photorealistic Animation and Rendering*. 2017, 1–10 1–7.
- [LMB*14] LIN, TSUNG-YI, MAIRE, MICHAEL, BELONGIE, SERGE, et al. "Microsoft coco: Common objects in context". *European conference on computer vision*. Springer. 2014, 740–755 4.
- [LZ21] LIU, SHIGUANG and ZHU, TING. "Structure-Guided Arbitrary Style Transfer for Artistic Image and Video". *IEEE Transactions on Multimedia* (2021) 6.
- [RBK21] RANFTL, RENÉ, BOCHKOVSKIY, ALEXEY, and KOLTUN, VLADLEN. "Vision Transformers for Dense Prediction". *ArXiv preprint* (2021) 2, 4, 6, 7.
- [RDB17] RUDER, MANUEL, DOSOVITSKIY, ALEXEY, and BROX, THOMAS. "Artistic style transfer for videos and spherical images". *CoRR* (2017). arXiv: [1708.04538](https://arxiv.org/abs/1708.04538). URL: <http://arxiv.org/abs/1708.04538> 1.
- [RLH*20] RANFTL, RENÉ, LASINGER, KATRIN, HAFNER, DAVID, et al. "Towards Robust Monocular Depth Estimation: Mixing Datasets for Zero-shot Cross-dataset Transfer". *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (2020) 2, 4, 6, 7.
- [SHKF12] SILBERMAN, NATHAN, HOIEM, DEREK, KOHLI, PUSHMEET, and FERGUS, ROB. "Indoor segmentation and support inference from rgb-d images". *European conference on computer vision*. Springer. 2012, 746–760 4.
- [SKLO18] SANAKOYEU, ARTSIOM, KOTOVENKO, DMYTRO, LANG, SABINE, and OMMER, BJORN. "A style-aware content loss for real-time hd style transfer". *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018, 698–714 1, 2.
- [SSN08] SAXENA, ASHUTOSH, SUN, MIN, and NG, ANDREW Y. "Make3d: Learning 3d scene structure from a single still image". *IEEE transactions on pattern analysis and machine intelligence* 31.5 (2008), 824–840 4.
- [SZ15] SIMONYAN, KAREN and ZISSERMAN, ANDREW. *Very Deep Convolutional Networks for Large-Scale Image Recognition*. 2015. arXiv: [1409.1556](https://arxiv.org/abs/1409.1556) [cs.CV] 2, 3.
- [ULVL16] ULYANOV, DMITRY, LEBEDEV, VADIM, VEDALDI, ANDREA, and LEMPITSKY, VICTOR S. "Texture networks: Feed-forward synthesis of textures and stylized images." *ICML*. Vol. 1. 2016, 4 1, 2.
- [UVL16] ULYANOV, DMITRY, VEDALDI, ANDREA, and LEMPITSKY, VICTOR. "Instance normalization: The missing ingredient for fast stylization". *arXiv preprint arXiv:1607.08022* (2016) 2, 3.
- [UVL17] ULYANOV, DMITRY, VEDALDI, ANDREA, and LEMPITSKY, VICTOR. "Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, 6924–6932 2.
- [YTBFI8] YEH, MAO-CHUANG, TANG, SHUAI, BHATTAD, ANAND, and FORSYTH, DAVID A. "Quantitative evaluation of style transfer". *arXiv preprint arXiv:1804.00118* (2018) 7.
- [ZMY21] ZHANG, JIAJING, MIAO, YONGWEL, and YU, JINHUI. "A Comprehensive Survey on Computational Aesthetic Evaluation of Visual Art Images: Metrics and Challenges". *IEEE Access* (2021) 7.