

# Deformable Style Transfer

Sunnie S. Y. Kim<sup>1</sup>[0000-0002-8901-7233], Nicholas Kolkin<sup>1</sup>[0000-0003-1233-1969],  
Jason Salavon<sup>2</sup>[0000-0003-4885-9511], and  
Gregory Shakhnarovich<sup>1</sup>[0000-0003-4700-9398]

<sup>1</sup>Toyota Technological Institute at Chicago  
{sunnie, nick.kolkin, greg}@ttic.edu

<sup>2</sup>University of Chicago  
salavon@uchicago.edu




Fig. 1: Each set of four images contains (left to right) a content input, a style input, a standard style transfer output, and our proposed method’s output. Sources of used images are available in the supplementary material.

**Abstract.** Both geometry and texture are fundamental aspects of visual style. Existing style transfer methods, however, primarily focus on texture, almost entirely ignoring geometry. We propose deformable style transfer (DST), an optimization-based approach that jointly stylizes the texture and geometry of a content image to better match a style image. Unlike previous geometry-aware stylization methods, our approach is neither restricted to a particular domain (such as human faces), nor does it require training sets of matching style/content pairs. We demonstrate our method on a diverse set of content and style images including portraits, animals, objects, scenes, and paintings. Code has been made publicly available at <https://github.com/sunniesuhyoung/DST>.

**Keywords:** Neural style transfer, geometric deformation, differentiable image warping

## 1 Introduction

The goal of style transfer algorithms is to re-render the content of one image using the style of one or several other images. Most modern approaches [7, 18,

2, 8, 14, 9, 20, 21, 17] capture a definition of “style” that focuses on color and texture. Art historians and other experts on image creation, however, define style more broadly and almost always consider the shapes and geometric forms present in an artwork as integral parts of its style [13, 6]. Shape and form play a vital role in recognizing the distinctive style of many artists in painting (e.g. Picasso, Modigliani, El Greco), sculpture (e.g., Botero, Giacometti), and other forms of media. While the results of style transfer algorithms thus far have been impressive and have captured the public’s attention, we propose a method of extending style transfer to better match the geometry of an artist’s style.

Style transfer methods that do not explicitly include geometry in their definition of style almost always keep the geometry of the content unchanged in the final output. This results in the outputs of these algorithms being easily identified as altered or “filtered” versions of the content image, rather than novel images created using the content image as a reference. Our focus in this work is to integrate shape and geometry as important markers of style and loosen the constraints on content as a receptive canvas. We achieve this by introducing a domain-agnostic geometric deformation of the content image, optimized jointly with standard style transfer losses.

Our proposed method, deformable style transfer (DST), takes two images as the input: a content image and a style image. We assume these two images share a domain and have some approximate alignment (e.g. both are images of sitting animals). This is a general scenario likely to arise in recreational or artistic uses of style transfer, as well as in tasks such as data augmentation. The nature of this problem makes *learning* to transfer style challenging since the variation in unconstrained domains and styles can not be reasonably presumed to be captured in any feasible training set. Therefore, like other style transfer work in this setting, we develop an optimization-based method, leveraging a pre-trained and fixed feature extractor derived from a convolutional network (CNN) trained for ImageNet classification.

There has been recent work on learning geometric style, using an explicit model of landmark constellations [25] or a deformation model representing a specific style [23]. These methods require a collection of images in the chosen style, and work only in a specific domain (often faces, due to their importance in culture and applications). Hence, they are not applicable to our more general scenario. Nonetheless, we compare our results to those of the aforementioned methods in their specific domain of faces in Section 5, and find that our method produces equally aesthetically pleasing results despite it being more general.

In this work we propose the first, to our knowledge, method for incorporating geometry into one-shot, domain-agnostic style transfer. The key idea of DST is to find a smooth deformation, or spatial warping, of the content image that brings it into spatial alignment with the style image. This deformation is guided by a set of matching keypoints, chosen to maximize the feature similarity between paired keypoints of the two images. After roughly aligning the paired keypoints with a rigid rotation and scaling, a simple  $\ell_2$  loss encourages warping our output image in such a way that the keypoints become spatially aligned. This

deformation loss is regularized with a total variation penalty to reduce artifacts due to drastic deformations, and combined with the more traditional style and content loss terms. DST’s joint, regularized objective simultaneously encourages preserving content, minimizing the style loss, and obtaining the desired deformation, weighing these goals against each other. This objective can be solved using standard iterative techniques.

To summarize the contributions of this work:

- We propose an optimization-based framework that endows style transfer algorithms with the explicit ability to deform a content image to match the geometry of a style image. Our flexible formulation also allows explicit user guidance and control of stylization tradeoffs.
- We demonstrate, for the first time, geometry-aware style transfer in a one-shot scenario. In contrast to previous works that are limited to human faces, DST works for images in other domains, with the assumption that they share a domain and have some approximate alignment.
- We evaluate DST on a range of style transfer instances, with images of faces, animals, vehicles, and landscapes, and through a user study demonstrate that our framework can augment existing style transfer algorithms to dramatically improve the perceived stylization quality, at minimal cost to the perceived content preservation.

## 2 Related Work

Early style transfer methods relied on hand-crafted features and algorithms [10–12, 5]. Gatys et al. [7] introduced Neural Style Transfer and dramatically improved the state-of-the-art by leveraging the features of a pretrained CNN. Neural Style Transfer represents “style” using the Gram matrix of features extracted from the shallow layers of a CNN and “content” using feature tensors extracted from the deeper layers. The pixels of a stylized output image are directly optimized to simultaneously match the style representation (of the style image) and content representation (of the content image). Subsequent works improve upon [7] with different complementary schemes, including spatial constraints [22], semantic guidance [3], and Markov Random Field priors [18]. Other work has improved upon [7] by replacing the objective function, style representation, and/or content representation [20, 9, 17].

Optimization-based methods produce high quality stylizations, but they can be computationally expensive as they require backpropagation at every iteration and gradually change the image, usually at a pixel level, until the desired statistics are matched. To overcome this limitation, model-based neural methods were introduced. These methods optimize a generative model offline, and at test time produce a stylized image with a single forward pass. These methods fall into two families with different tradeoffs relative to optimization-based style transfer. Some methods [21] trade flexibility for speed and quality, quickly producing excellent stylizations but only for a predetermined set of styles. Other methods [15, 14] trade off quality for speed, allowing for fast transfer of arbitrary

styles, but typically produce lower quality outputs than optimization-based style transfer. Each family of method excels in a different regime, and in this work we prioritize flexibility and quality over speed.

Until recently, style transfer methods could not transfer geometric style and were limited to transferring color and texture. In much work outside the domain of style transfer, however, geometric transformation has been applied to images via automatic warping. Early work required predicting a set of global transformation parameters or a dense deformation field. Cole et al. [4] enabled fine-grained local warping by proposing a method that takes a sparse set of control keypoints and warps an image with spline interpolation. The introduced warping module is differentiable and thus can be trained as part of an end-to-end system, although [4] requires pre-detected landmarks as input for its face synthesis task.

Several recent works have attempted to combine image warping with neural networks to learn both textural and geometric style of human portraits. CariGAN [19] translates a photo to a caricature by training a Generative Adversarial Network (GAN) that models geometric transformation with manually annotated facial landmarks and another GAN that translates the usual non-geometric style appearances. Face of Art (FoA) [25] trains a neural network model to automatically detect 68 canonical facial landmarks in artistic portraits and uses them to warp a photo so that the geometry of the face is closer to that of an artistic portrait. WarpGAN [23], on the other hand, adds a warping module to the generator of a GAN and trains it as part of an end-to-end system by optimizing both the locations of keypoints and their displacements, using a dataset that contains a caricature and photo pair for each identity.

The main distinction of our work from these efforts is that ours is not limited to human faces (or any other particular domain) and does not require offline training on a specially prepared dataset. In terms of methodology, FoA separates transferring “texture” and transferring geometry, while DST transfers them jointly. WarpGAN treats texture and geometry jointly, but it has to learn a warping module with paired examples in the face caricature domain while DST does not. We show in Section 5 that results of our more general method, even when applied to faces, are competitive with results of these two face-specific methods.


For finding correspondences between images, two recent efforts use CNN-based descriptors to identify matching points between paired images outside the domain of human faces. Fully Convolutional Self-Similarity [16] is a descriptor for dense semantic correspondence that uses local self-similarity to match keypoints among different instances within the same object class. Neural Best-Buddies (NBB) [1] is a more general method for finding a set of sparse cross-domain correspondences that leverages the hierarchical encoding of features by pre-trained CNNs. We use NBB in our method with post-processing, as described in detail in the next section.



### 3 Geometry Transfer via Correspondences

One path for introducing geometric style transfer is establishing spatial associations between the content and style images, and defining a *deformation* that brings the content image into (approximate) alignment with the style image. Assuming they share a domain and have similar geometry (e.g. both are images of front-facing cars), we can aim to find meaningful spatial correspondences to define the deformation. The correspondences would specify displacement “targets”, derived from the style image, for keypoints in the content image. Thin-plate spline interpolation [4] can extend this sparse set of displacements to a full displacement field specifying how to deform every pixel in the output image.

Fig. 2: Illustration of our method using keypoints taken from FoA (rows 1-3) or generated manually (row 4). Keypoints are overlaid on the content and style images with matching points in the same color. Naive warp indicates output of standard style transfer warped by moving source points to target points. Figure is best viewed zoomed-in on screen.



#### 3.1 Finding and Cleaning Keypoints

If we fix a domain and assume availability of a training set drawn from the domain, we may be able to learn a domain-specific mechanism for finding salient and meaningful correspondences. This can be done through facial landmark detection [25] or through learning a data-driven detector for relevant points [16, 23]. Alternatively, we can expect a user interacting with a style transfer tool to manually select points they consider matching in the two images. If matching points are provided by such approaches, they can be used in DST as we show in Figure 2. However, we are interested in a more general scenario, a one-shot, domain-agnostic setting where we may not have access to such points. Hence we use NBB, a generic method for point matching between images.

NBB finds a sparse set of correspondences between two images that could be from different domains or semantic categories. It utilizes the hierarchy of deep features of a pre-trained CNN, i.e. the characteristic that deeper layers extract

high-level, semantically meaningful, and spatially invariant features and shallow layers encode low-level features such as edges and color features. Starting from the deepest layer, NBB searches for pairs of correspondences that are mutual nearest neighbors, filters them based on activation values, and percolates them through the hierarchy to narrow down the search region at each level. At the end of the algorithm, it clusters the set of pixel-level correspondences into  $k$  spatial clusters and returns  $k$  keypoint pairs.

The keypoint pairs returned by NBB, however, are often too noisy and not sufficiently spread out to use them as they are. To provide better guidance for geometric deformation, we modify NBB to get a cleaner and better spatially-distributed set of pairs. Specifically, we remove the final clustering step and return all pixel-level correspondences, usually on the order of hundreds of correspondence pairs. Then we use a greedy algorithm that selects a keypoint with the highest activation value (calculated by NBB) that is at least 10 pixels away from any already selected keypoint. We select up to 80 keypoint pairs and filter out keypoints with activation values smaller than 1. After the initial selection, we map the keypoints in the style image onto the content image by finding a similarity transformation that minimizes the squared distance between the two point clusters [24]. We then additionally clean up the selected keypoints by removing keypoints pairs that cross each other, to prevent a discontinuous warp field. (If keypoints are provided by FoA, manual selection, or other non-NBB methods, we skip the cleaning process and only transform the style image keypoints appropriately.) We refer to the keypoints in the content image as the “source points” and the corresponding keypoints in the style image mapped onto the content image as the “target points.” This process is illustrated in Figure 3.

### 3.2 Differentiable Image Warping

We specify an image deformation by a set of source keypoints  $P = \{p_1, \dots, p_k\}$  and the associated 2D displacement vectors  $\theta = \{\theta_1, \dots, \theta_k\}$ .  $\theta$  specify for each source keypoint  $p_i$ , the *destination* coordinates  $p_i + \theta_i$ . Following [23] we use thin-plate spline interpolation [4] to produce a dense flow field from the coordinates of an unwarped image  $I$  to a warped image  $W(I, \theta)$ . This is a closed-form procedure which finds parameters  $w, v, b$  that minimize  $\sum_{i=1}^k \|f_\theta(p_i + \theta_i) - p_i\|^2$  subject to a curvature constraint. With these parameters, we have the inverse mapping function

$$f_\theta(q) = \sum_{i=1}^k w_i \phi(\|q - p_i - \theta_i\|) + v^T q + b \quad (1)$$

where  $q$  denotes the location of a pixel in the warped image and  $\phi$  is a kernel function which we choose to be  $\phi(r) = r^2 \log(r)$ .  $f_\theta(q)$  gives the inverse mapping of the pixel  $q$  in the original image; i.e. the pixel coordinates in the unwarped image from which we should derive the color of pixel  $q$  in the warped image. The color of each pixel can then be generated through bilinear sampling. This entire warping module is differentiable with respect to  $\theta$ , allowing it to be optimized as part of an end-to-end system.

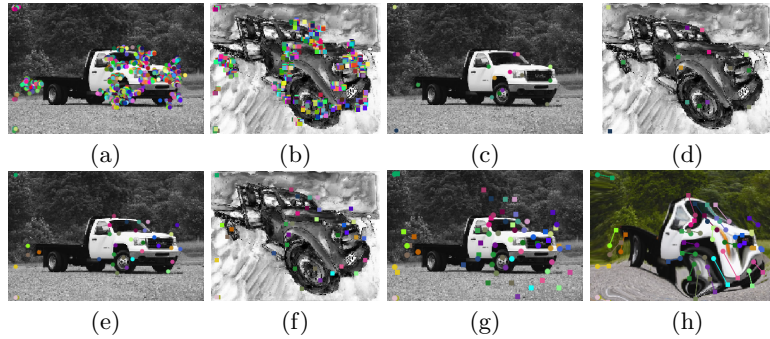


Fig. 3: An image can be spatially deformed by moving a set of source points to a set of target points. Matching keypoints are indicated by color. (a) Content image with all correspondences. (b) Style image with all correspondences. (c) Content image with original NBB keypoints. (d) Style image with original NBB keypoints. (e) Content image with our selected keypoints. (f) Style image with our selected keypoints. (g) Content image with keypoints aligned by just matching the centers. (h) Content image warped with keypoints aligned by a similarity transformation. The lines indicate where the circle source points move to (square target points). Figure is best viewed zoomed-in on screen.

## 4 Spatially Guided Style Transfer

The input to DST consists of a content image  $I_c$ , a style image  $I_s$ , and aligned keypoint pairs  $P$  (source) and  $P'$  (target). Recall that these points don't have to be infused with explicit domain- or category-specific semantics. DST optimizes the stylization parameters (usually the pixels of the output image)  $X$  and the deformation parameters  $\theta$ . The final output is the warped stylized image  $W(X, \theta)$ .

### 4.1 Content and Style Loss Terms

DST can be used with any one-shot, optimization-based style transfer method with a content loss and a style loss. In this work, we demonstrate our framework with two such methods: Gatys et al. [7] and Kolkin et al. [17], which we will refer to as Gatys and STROTSS, respectively. Each method defines a content loss  $L_{\text{content}}(I_c, X)$  and a style loss  $L_{\text{style}}(I_s, X)$ . These aim to capture the visual content of  $I_c$  and the visual style of  $I_s$  in the output  $X$ . Below we briefly summarize the content/style loss of these methods. For more details, we direct the reader to [7, 17].

Gatys represents “style” in terms of the Gram matrix of features extracted from multiple layers of a CNN and “content” as the feature tensors extracted from another set of layers. It defines  $L_{\text{content}}$  as the squared-error between the feature representation of the content image and that of the output image. Sim-

ilarly, it defines  $L_{\text{style}}$  as the weighted sum of the squared-error between the Gram matrix of the style image and that of the output image.

STROTSS, inspired by the concept of self-similarity, defines  $L_{\text{content}}$  as the absolute error between the normalized pairwise cosine distance between feature vectors extracted from the content image and those of the output image. STROTSS's  $L_{\text{style}}$  is composed of three terms: the Relaxed Earth Movers Distance (EMD), the moment matching term, and the color matching term. Relaxed EMD helps transfer the structural forms of the style image to the output image. The moment matching term, which aims to match the mean and covariance of the feature vectors in the two images, combats over- or under-saturation. The color matching term, defined as the Relaxed EMD between pixel colors, encourages the output image and the style image to have a similar palette.

When using DST with a base style transfer method, we do not change anything about  $L_{\text{content}}$ . The style loss of DST is composed of two terms

$$L_{\text{style}}(I_s, X) + L_{\text{style}}(I_s, W(X, \theta)). \quad (2)$$

The first loss term is between the style image  $I_s$  and the *unwarped* stylized image  $X$ . The second term is between  $I_s$  and the spatially deformed stylized image  $W(X, \theta)$ , with  $\theta$  defining the deformation as per Section 3. Minimizing Eq. (2) is aimed at finding a good stylization both with and without spatial deformation. This way we force the stylization parameters  $X$  and the spatial deformation parameters  $\theta$  to work together to produce a harmoniously stylized and spatially deformed final output  $W(X, \theta)$ .

## 4.2 Deformation Loss Term

Given a set of  $k$  source points  $P$  and matching target points  $P'$ , we define the deformation loss as

$$L_{\text{warp}}(P, P', \theta) = \frac{1}{k} \sum_{i=1}^k \|p'_i - (p_i + \theta_i)\|_2, \quad (3)$$

where  $p_i$  and  $p'_i$  are the  $i$ -th source and target point coordinates. Minimizing Eq. (3) with respect to  $\theta$  seeks a set of displacements that move the source points to the target points. This term encourages the geometric shape of the stylized image to become closer to that of the style.

Aggressively minimizing the deformation loss may lead to significant artifacts, due to keypoint match errors or incompatibility of the content with the style geometry. To avoid these artifacts, we add a regularization term encouraging smooth deformations. Specifically, we use the (anisotropic) total variation norm of the 2D warp field  $f$  normalized by its size

$$R_{\text{TV}}(f) = \frac{1}{W \times H} \sum_{i=1}^W \sum_{j=1}^H \|f_{i+1,j} - f_{i,j}\|_1 + \|f_{i,j+1} - f_{i,j}\|_1. \quad (4)$$

This regularization term smooths the warp field by encouraging nearby pixels to move in a similar direction.

### 4.3 Joint Optimization

Putting everything together, the objective function of DST is

$$\begin{aligned} L(X, \theta, I_c, I_s, P, P') = & \alpha L_{\text{content}}(I_c, X) \\ & + L_{\text{style}}(I_s, X) + L_{\text{style}}(I_s, W(X, \theta)) \\ & + \beta L_{\text{warp}}(P, P', \theta) \\ & + \gamma R_{\text{TV}}(f_\theta), \end{aligned} \quad (5)$$

where  $X$  is the stylized image and  $\theta$  parameterizes the spatial deformation. Hyperparameters  $\alpha$  and  $\beta$  control the relative importance of content preservation and spatial deformation to stylization. Hyperparameter  $\gamma$  controls the amount of regularization on the spatial deformation. The effect of varying  $\alpha$  is analyzed in [7, 17]. The effect of changing  $\beta$  and  $\gamma$  is illustrated in Figure 4. We use standard iterative techniques such as stochastic gradient descent or L-BFGS to minimize Eq. (5) with respect to  $X$  and  $\theta$ . Implementation details can be found in the supplementary material and the published code.<sup>1</sup>




Fig. 4: DST outputs with varying  $\beta$  and  $\gamma$  using STROTSS as the base method. Image in the upper right corner (low  $\beta$ , high  $\gamma$ ) has the least deformation, and the image in the bottom left corner (high  $\beta$ , low  $\gamma$ ) has the most deformation.

## 5 Results

We observe that DST often captures the geometric style of the target style image. One visually striking effect is that the resulting images no longer look like “filtered” versions of the original content, as they often do with standard style transfer methods. We show results of DST with Gatys and STROTSS in Figures 5 and 6. For a pair of content and style images, we show the output of

<sup>1</sup> <https://github.com/sunniesuhyoung/DST>

DST and the output of Gatys/STROTSS that doesn't have the spatial deformation capability. To highlight the effect of the DST-learned deformation, we also provide the content image warped by DST and the Gatys/STROTSS output naively warped with the selected keypoints without any optimization of the deformation. While naive warping produces undesirable artifacts, DST finds a warp that harmoniously improves stylization while preserving content.

As a simple quantitative evaluation, we calculated the (STROTSS) style loss on 185 pairs of DST and STROTSS outputs. Surprisingly, we found that on average this loss was 7% higher for DST outputs than STROTSS ones, even for examples we show in Figure 6. While the loss difference is small, this is a mismatch with the human judgment of stylization quality shown in Section 5.2.




Fig. 5: DST results with Gatys. \*Naively warped by moving source points to target points. †Warp learned by DST applied to the content image.

### 5.1 Comparison with FoA and WarpGAN

While, so far as we are aware, DST is the first work to allow open-domain geometry-aware style transfer, other work has tackled this problem in the domain of human faces. To compare performance, we show results of DST and results of FoA [25] and WarpGAN [23] on the same content-style pairs in Figure 7. Note that both of these methods require training a model on a dataset of stylized portraits or caricatures, while DST operates with access to only a single content and single style image.





Fig. 6: DST results with STROTSS. \*Naively warped by moving source points to target points. †Warp learned by DST applied to the content image.




Fig. 7: Comparison of DST with (a) FoA and (b) WarpGAN. †: Note that WarpGAN’s does not use a specific style image, so this style image is only used by DST; see text for details.

DST jointly optimizes the geometric and non-geometric stylization parameters, while FoA transfers geometric style by warping the facial landmarks in the content image to a specific artist’s (e.g. Modigliani) canonical facial landmark pattern (with small variations added) learned by training a model on a dataset of stylized portraits. FoA then separately transfers textural style with a standard style transfer method (e.g. Gatys, STROTSS). When we compare DST and FoA in Figure 7, we demonstrate “one-shot FoA” since the style images used to produce the outputs in [25] are unavailable. That is, we assume that we have access to one content image, one style image, and the trained FoA landmark detector. Using the detector, we find 68 facial landmarks in the content and style images and transform the style image landmarks, as described in Section 3, to get the target points. Then we follow FoA’s two-step style transfer and transfer the textural style by STROTSS and transfer the geometric style by warping the output image by moving the source points to the target points.

The biggest difference between WarpGAN and DST is that DST is a one-shot style transfer method that works with a single content image and a single style image. WarpGAN, on the other hand, is trained on a dataset of paired pictures and caricatures of the same identities, and generates a caricature for an input content image from its learned deformation model. To compare the performance of WarpGAN and DST, we used content/style image pairs from [23] and ran DST. In Figure 7, we show the outputs of DST and the outputs of WarpGAN taken from [23]. Despite the lack of a *learning* component, DST results are competitive and sometimes more aesthetically pleasing than results of FoA and WarpGAN.

## 5.2 Human Evaluation

Quantitatively evaluating and comparing style transfer is challenging, in part because of the subjective nature of aesthetic properties defining style and visual quality, and in part due to the inherent tradeoff between content preservation and stylization [26, 17]. Following the intuition developed in these papers, we conducted a human evaluation study using Amazon Mechanical Turk on a set of 75 diverse style/content pairs. The goal was to study the effect of DST on the stylization/content preservation tradeoff, in comparison to the base style transfer methods. The evaluation was conducted separately for STROTSS and Gatys-based methods. We considered three DST deformation regimes: low ( $\beta = 0.3, \gamma = 75$ ), medium ( $\beta = 0.5, \gamma = 50$ ), and high ( $\beta = 0.7, \gamma = 10$ ) for STROTSS; low ( $\beta = 3, \gamma = 750$ ), medium ( $\beta = 7, \gamma = 100$ ), and high ( $\beta = 15, \gamma = 100$ ) for Gatys. So for each base method, we compare four stylized output images. The effect of varying  $\beta$  and  $\gamma$  is illustrated in Figure 4.

To measure content preservation, we asked MTurk users the question: “Does image A represent the same scene as image B”, where A referred to the content image and B to the output of style transfer. The users were forced to choose one of four answers: “Yes”, “Yes, with minor errors”, “Yes, with major errors” and “No”. Converting these answers to numerical scores (1 for “No”, 4 for “Yes”) and



averaging across content/style pairs and users, we get a *content score* between 1 and 4 for each of the four methods.

To evaluate the effect of the proposed deformable framework, we presented the users with a pair of outputs, one from the base method (Gatys or STROTSS) and the other from DST, along with the style image. The order of the first two is randomized. We asked the users to choose which of the two output images better matches the style. The fraction of time a method is preferred in all comparisons (across methods compared to, users, content/style pairs) gives a *style score* between 0 and 1. 0.7 means that the method “wins” 70% of all comparisons it was a part of. The evaluation interfaces are provided in the supplementary material.

In total, there were 600 unique content comparisons: 4 questions  $\times$  75 images for Gatys and an equal number for STROTSS. 123 users participated in the evaluation, and each comparison was evaluated by 9.55 users on average. The standard deviation of the content choice agreement was 0.79 (over a range of 1 to 4). For stylization, there were 450 unique comparisons in total: 3 comparisons between the base method and each of the 3 DST deformation regimes  $\times$  75 images for Gatys and likewise for STROTSS. 103 users participated in the stylization evaluation, and each comparison was evaluated by 8.76 users on average. For each comparison, 6.47 users agreed in their choice on average.

Results of this human evaluation are shown in Figure 8. Across the deformation regimes (low, medium, high), for both STROTSS and Gatys, DST significantly increases the perceived stylization quality, while only minimally reducing the perceived content preservation. Note that some reduction to the content score can be expected since we intentionally alter the content more by deforming it, but our evaluation shows that this drop is small.




Fig. 8: Human evaluation results, comparing DST in different deformation regimes with STROTSS (green) and Gatys (blue). DST provides a much higher perceived degree of style capture without a significant sacrifice in content preservation.

### 5.3 Limitations

In Figure 9, we show unsuccessful examples of DST where the output image did not deform towards having a similar shape as the style image or deformed only partially. We observed that bad deformations often stem from poorly matching or too sparse set of keypoints. We expect finding better matching keypoints between images and making the method more robust to poor matches will improve results.




Fig. 9: Examples of DST failures. We observed that stylization failures are often due to correspondence errors or overly complex scene layout.

## 6 Conclusion

Prior work on style transfer has largely ignored geometry and shape, despite the role these play in visual style. We present deformable style transfer (DST), a novel approach that combines the traditional texture and color transfer with spatial deformations. Our method incorporates deformation targets, derived from domain-agnostic point matching between content and style images, into the objective of an optimization-based style transfer framework. This is to our knowledge the first effort to develop a one-shot, domain-agnostic method for capturing and transferring geometric aspects of style.

Still many aspects of geometric style transfer remain unexplored. Narrowing this might involve developing more robust keypoint matching algorithms for highly stylized images. Furthermore, it is by no means certain that modifying geometry using warp fields driven by paired keypoints is the most effective approach to this problem. We hope that future work will continue to explore how to represent geometric style more flexibly, and more accurately encode the artistic extraction and abstraction of shape and form.

## Acknowledgement

This work is supported by the DARPA GARD award HR00112020003 and the University of Chicago CDAC Discovery Grant. We would also like to thank Haochen Wang, Davis Gilton, Steven Basart, and members of the Perception and Learning Systems (PALS) at TTIC for their helpful comments.

## References

1. Aberman, K., Liao, J., Shi, M., Lischinski, D., Chen, B., Cohen-Or, D.: Neural best-buddies: Sparse cross-domain correspondence. *ACM Transactions on Graphics (TOG)* **37**(4), 69 (2018)
2. Berger, G., Memisevic, R.: Incorporating long-range consistency in cnn-based texture generation. *arXiv preprint arXiv:1606.01286* (2016)
3. Champandard, A.J.: Semantic style transfer and turning two-bit doodles into fine artworks. *CoRR* **abs/1603.01768** (2016), <http://arxiv.org/abs/1603.01768>
4. Cole, F., Belanger, D., Krishnan, D., Sarna, A., Mosseri, I., Freeman, W.T.: Synthesizing normalized faces from facial identity features. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (July 2017)
5. Efros, A.A., Freeman, W.T.: Image quilting for texture synthesis and transfer. In: *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. pp. 341–346. *ACM* (2001)
6. Elkins, J.: Style. *The Dictionary of art* **29** (1996)
7. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016)
8. Gatys, L.A., Ecker, A.S., Bethge, M., Hertzmann, A., Shechtman, E.: Controlling perceptual factors in neural style transfer. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
9. Gu, S., Chen, C., Liao, J., Yuan, L.: Arbitrary style transfer with deep feature reshuffle. In: *CVPR* (2018)
10. Haeberli, P.: Paint by numbers: Abstract image representations. In: *ACM SIGGRAPH computer graphics*. vol. 24, pp. 207–214. *ACM* (1990)
11. Hertzmann, A.: Painterly rendering with curved brush strokes of multiple sizes. In: *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. pp. 453–460. *ACM* (1998)
12. Hertzmann, A., Jacobs, C.E., Oliver, N., Curless, B., Salesin, D.H.: Image analogies. In: *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. pp. 327–340. *ACM* (2001)
13. Hofstadter, D.: Metamagical theamas: Variations on a theme as the essence of imagination. *Scientific American* **247**(4) (1983)
14. Huang, X., Belongie, S.J.: Arbitrary style transfer in real-time with adaptive instance normalization. In: *ICCV* (2017)
15. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *European Conference on Computer Vision*. pp. 694–711. *Springer* (2016)
16. Kim, S., Min, D., Ham, B., Lin, S., Sohn, K.: Fcsc: Fully convolutional self-similarity for dense semantic correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(3), 581–595 (March 2019). <https://doi.org/10.1109/TPAMI.2018.2803169>
17. Kolkin, N., Salavon, J., Shakhnarovich, G.: Style transfer by relaxed optimal transport and self-similarity. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
18. Li, C., Wand, M.: Combining markov random fields and convolutional neural networks for image synthesis. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016)

19. Li, W., Xiong, W., Liao, H., Huo, J., Gao, Y., Luo, J.: Carigan: Caricature generation through weakly paired adversarial learning. CoRR **abs/1811.00445** (2018), <http://arxiv.org/abs/1811.00445>
20. Mechrez, R., Talmi, I., Zelnik-Manor, L.: The contextual loss for image transformation with non-aligned data. arXiv preprint arXiv:1803.02077 (2018)
21. Sanakoyeu, A., Kotovenko, D., Lang, S., Ommer, B.: A style-aware content loss for real-time hd style transfer. In: CVPR (2018)
22. Selim, A., Elgharib, M., Doyle, L.: Painting style transfer for head portraits using convolutional neural networks. ACM Trans. Graph. **35**(4), 129:1–129:18 (Jul 2016). <https://doi.org/10.1145/2897824.2925968>, <http://doi.acm.org/10.1145/2897824.2925968>
23. Shi, Y., Deb, D., Jain, A.K.: Warpgan: Automatic caricature generation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
24. Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. IEEE Transactions on Pattern Analysis and Machine Intelligence **13**(4), 376–380 (April 1991). <https://doi.org/10.1109/34.88573>
25. Yaniv, J., Newman, Y., Shamir, A.: The face of art: Landmark detection and geometric style in portraits. ACM Trans. Graph. **38**(4), 60:1–60:15 (Jul 2019). <https://doi.org/10.1145/3306346.3322984>, <http://doi.acm.org/10.1145/3306346.3322984>
26. Yeh, M.C., Tang, S., Bhattad, A., Forsyth, D.A.: Quantitative evaluation of style transfer. arXiv preprint arXiv:1804.00118 (2018)