Joint Bilateral Learning for Real-time Universal **Photorealistic Style Transfer**

Xide Xia^{1*}, Meng Zhang^{2*}, Tianfan Xue³, Zheng Sun³, Hui Fang^{3*}, Brian Kulis¹, and Jiawen Chen³

¹ Boston University {xidexia,bkulis}@bu.edu ² PixelShift.AI meng@pixelshift.ai ³ Google Research {tianfan,zhengs,hfang,jiawen}@google.com

(a) Input content image

(b) Style 1







(c) Style 2

Fig. 1: Our method takes an input image (a) and renders it in the style of an arbitrary reference photo (insets) while preserving scene content. Notice that although the stylized outputs (b) and (c) differ dramatically in appearance from the input, one property they share is that nearby pixels of similar color are transformed similarly. We visualize this in the grayscale intensity domain, where the transformation is approximately a curve (zoomed insets). Taking advantage of this property, we propose a feed-forward neural network that directly learns these local curves.

Abstract. Photorealistic style transfer is the task of transferring the artistic style of an image onto a content target, producing a result that is plausibly taken with a camera. Recent approaches, based on deep neural networks, produce impressive results but are either too slow to run at practical resolutions, or still contain objectionable artifacts. We propose a new end-to-end model for photorealistic style transfer that is both fast and inherently generates photorealistic results. The core of our approach is a feed-forward neural network that learns local edge-aware affine transforms that automatically obey the photorealism constraint. When trained on a diverse set of images and a variety of styles, our model can robustly apply style transfer to an arbitrary pair of input images. Compared to the state of the art, our method produces visually superior results and is three orders of magnitude faster, enabling realtime performance at 4K on a mobile phone. We validate our method with ablation and user studies.

Keywords: style transfer, bilateral learning, local affine transform

^{*} Work done at Google Research.

1 Introduction

Image style transfer has recently received significant attention in the computer vision and machine learning communities [9]. A central problem in this domain is the task of transferring the style of an arbitrary image onto a *photorealistic* target. The seminal work of Gatys et al. [4] formulates this general artistic style transfer problem as an optimization that minimizes both style and content losses. but results often contain spatial distortion artifacts. Luan et al. [20] seek to reduce these artifacts by adding a *photorealism constraint*, which encourages the transformation between input and output to be *locally affine*. However, because the method formulates the problem as a large optimization whereby the loss over a deep network must be minimized for every new image pair, performance is limited. The recent work of Yoo et al. [32] proposes a wavelet-based method which provides stable stylization but is not fast enough to run at practical resolutions. Another line of recent work seeks to pretrain a feed-forward deep model [3, 8, 10, 14, 16, 18, 29, 30] that once trained, can produce a stylized result with a single forward pass at test time. While these "universal" [9] techniques are significantly faster than those based on optimization, they may not generalize well to unseen images, may produce non-photorealistic results, and are still to slow to run in real-time on a mobile device.

In this work, we introduce a fast end-to-end method for photorealistic style transfer. Our model is a single feed-forward deep neural network that once trained on a suitable dataset, runs in real-time on a mobile phone at full camera resolution (i.e., 12 megapixels or "4K")—significantly faster than the state of the art. Our key observation is that we can guarantee photorealistic results by strictly enforcing Luan et. al's photorealism constraint [20]—locally, regions of similar color in the input must map to a similarly colored region in the output while respecting edges. Therefore, we design a deep learning algorithm in *bilateral space*, where these local affine transforms can be compactly represented. We contribute:

- 1. A photorealistic style transfer network that learns local affine transforms. The model is robust and degrades gracefully when confronted with unseen or adversarial inputs.
- 2. An inference implementation that runs in real-time at 4K on a mobile phone.
- 3. A bilateral-space Laplacian regularizer eliminates spatial grid artifacts.

1.1 Related Work

Early work in image style transfer can be classified as either global methods that transfer statistics over the entire image [22, 24] or local methods [13, 25, 26, 31, 28] that typically find dense correspondences between content and style. While global methods are efficient, the results are not always faithful to the target style. Local methods can generate high-quality results, but they are computationally expensive and often fail to capture semantics. We highlight that these techniques

3



Fig. 2: **Inspiration.** Artistic style transfer methods such as AdaIN generalize well to diverse content/style inputs but exhibit distortions on photographic content. HDRnet, designed to reproduce arbitrary imaging operators, learns the desired transform representation but fails to capture universal style transfer. Our work combines ideas from both approaches.

do produce photorealistic results, albeit not always faithful to the style or well exposed.

Recently, Gatys et al. [4] showed that style can be effectively captured by the statistics of layer activations within deep neural networks trained for discriminative image classification. However, due to its generality, the technique and its successors often contain non-photorealistic painterly spatial distortions. To remove such distortions, He et al. [7] propose to achieve a more accurate color transfer by leveraging semantically-meaningful dense correspondence between images. Luan et al. [20] ameliorate this problem by imposing additional constraints on the loss function. The result is pushed towards photorealism by constraining the transformation to be locally affine in color space.

PhotoWCT [18] imposes a similar constraint as a postprocessing step, while LST [14] appends a spatial propagation network [19] after the main style transfer network to preserves the desired affinity. Similarly, Puy et al. [23] propose a flexible network to perform artistic style transfer, and applies postprocessing after each learned update for photorealistic content. Compared to these ad hoc approaches where the photorealism constraint is a soft penalty, our model directly predicts local affine transforms, guaranteeing that the constraint is satisfied.

Another line of recent work shows that matching the statistics of autoencoders is an effective way to parameterize style transfer [8, 17, 18, 14, 32]. Moreover, they show that distortions can be reduced by preserving high frequencies using unpooling [18] or wavelet transform residuals [32].

Our work unifies these two lines of research. Our network architecture builds upon HDRnet [5], which was first employed in the context of learning image enhancement and tone manipulation. Given a large dataset of input/output pairs, it learns local affine transforms that best reproduces the operator. The network is small and the learned transforms that are intentionally constrained to be incapable of introducing artifacts such as noise or false edges. These are exactly the properties we want and indeed, Gharbi et. al. demonstrated style transfer in their original paper. However, when we applied HDRnet to our more diverse dataset, we found a number of artifacts (Figure 2). This is because HDRnet does

not explicitly model style transfer and instead learns by memorizing what it sees during training and projecting the function onto local affine transforms. Therefore, it will require a lot of training data and generalize poorly. Since HDRnet learns local affine transforms from low-level image features, our strategy is to start with statistical feature matching using Adaptive Instance Normalization [8] to build a joint distribution. By explicitly modeling the style transformation as a distribution matching process, our network is capable of generalizing to unseen or adversarial inputs (Figure 2).

2 Method

Our method is based on a single feed-forward deep neural network. It takes as input two images, a content photo I_c , and an arbitrary style image I_s , producing a photorealistic output O with the former's content but the latter's style. Similar to previous work [18, 14, 32], our network is "universal"—after training on a diverse dataset of content/style pairs, it can generalize to novel input combinations. Its architecture is centered around the core idea of learning local affine transformations, which inherently enforce the photorealism constraint.

2.1 Background

For completeness, we first summarize the key ideas of recent work.

Content and Style. The Neural Style Transfer [4] algorithm is based on an optimization that minimizes a loss balancing the output image's fidelity to the input images' content and style:

$$\mathcal{L}_g = \alpha \mathcal{L}_c + \beta \mathcal{L}_s \qquad \text{with} \qquad (1)$$

$$\mathcal{L}_{c} = \sum_{i=1}^{N_{c}} \|F_{i}[O] - F_{i}[I_{c}]\|_{2}^{2} \text{ and } \mathcal{L}_{s} = \sum_{i=1}^{N_{s}} \|G_{i}[O] - G_{i}[I_{s}]\|_{F}^{2}, \quad (2)$$

where N_c and N_s denote the number of intermediate layers selected from a pretrained VGG-19 network [27] to represent image content and style, respectively. Scene content is captured by the feature maps F_i of intermediate layers of the VGG network, and style is captured by their Gram matrices $G_i[\cdot] = F_i[\cdot]F_i[\cdot]^T$. $||\cdot||_F$ denotes the Frobenius norm.

Statistical Feature Matching. Instead of directly minimizing the loss in Equation 1, followup work shows that it is more effective to match the statistics of feature maps at the bottleneck of an auto-encoder. Variants of the whitening and coloring transform [17, 18, 32] normalize the singular values of each channel, while Adaptive Instance Normalization (AdaIN) [8] proposes a simple scheme using the mean $\mu(\cdot)$ and the standard deviation $\sigma(\cdot)$ of each channel:

AdaIN
$$(x, y) = \sigma(y) \left(\frac{x - \mu(x)}{\sigma(x)}\right) + \mu(y),$$
 (3)



Fig. 3: **Model architecture.** Our model starts with a low-resolution grid prediction stream that uses three feature transfer blocks S to build a joint distribution between the low-level features of the input content/style pair. This distribution is fed to bilateral learning blocks L and G to predict an affine bilateral grid Γ . Rendering, which runs at full-resolution, performs the minimal per-pixel work of sampling from Γ a 3 × 4 matrix and then multiplying.

where x and y are content and style feature channels, respectively. Due to its simplicity and reduced cost, we also adopt AdaIN layers in our network architecture as well as its induced style loss [8, 15]:

$$\mathcal{L}_{sa} = \sum_{i=1}^{N_S} \|\mu(F_i[O]) - \mu(F_i[I_s])\|_2^2 + \sum_{i=1}^{N_S} \|\sigma(F_i[O]) - \sigma(F_i[I_s])\|_2^2.$$
(4)

Bilateral Space. Bilateral space was first introduced by Paris and Durand [21] in the context of fast edge-aware image filtering. A 2D grayscale image I(x, y) can be "lifted" into bilateral space as a sparse collection of 3D points $\{x_j, y_j, I_j\}$ in the augmented space. In this space, linear operations are inherently edge-aware because Euclidean distances preserve edges. They prove that bilateral filtering is equivalent to *splatting* the input onto a regular 3D *bilateral grid*, blurring, and *slicing* out the result using trilinear interpolation at the input coordinates $\{x_j, y_j, I_j\}$. Since blurring and slicing are low-frequency operations, the grid can be low-resolution, dramatically accelerating the filter.

Bilateral Guided Upsampling (BGU) [2] extends the bilateral grid to represent transformations between images. By storing at each cell an affine transformation, an *affine bilateral grid* can encode any image-to-image transformation given sufficient resolution. The pipeline is similar: *splat* both input and output images onto a bilateral grid, blur, and perform a per-pixel least squares fit. To apply the transform, *slice* out a per-pixel affine matrix and multiply by the input color. BGU shows that this representation can accelerate a variety of imaging operators and that the approximation degrades gracefully with resolution when suitably regularized. Affine bilateral grids are constrained to produce an output

	S_1^1	S_1^2	S_2^1	S_2^2	S_3^1	S_3^2	C^7	C^8	L^1	L^2	G^1	G^2	G^3	G^4	G^5	G^6	F	Г
type	c	c	c	c	c	c	c	c	c	c	c	c	f_c	f_c	f_c	f_c	c	c
stride	2	1	2	1	2	1	2	1	1	1	2	2	-	-	-	-	1	1
size	128	128	64	64	32	32	16	16	16	16	8	4	-	-	-	-	16	16
channels	8	8	16	16	32	32	64	64	64	64	64	64	256	128	64	64	64	96

Table 1: Details of our network architecture. S_j^i denotes the *i*-th layer in the *j*-th transfer block. We apply AdaIN after each S_j^1 . L^i , G^i , F, and Γ refer to local features, global features, fusion, and learned bilateral grid, respectively. Local and global features are concatenated before fusion F. c and f_c denote convolutional and fully-connected layers, respectively. Convolutions are all 3×3 except F, where it is 1×1 .

that is a smoothly varying, edge-ware, and locally affine transformation of the input. Therefore, it fundamentally cannot produce false edges, amplify noise, and inherently obeys the photorealism constraint.

Gharbi et al. [5] showed that slicing and applying an affine bilateral grid are sub-differentiable and therefore can be incorporated as a layer in a deep neural network and learned using gradient descent. They demonstrated that their HDRnet architecture can effectively learn to reproduce many photographic tone mapping and detail manipulation tasks, regardless of whether they are algorithmic or artist-driven.

2.2 Network Architecture

Our end-to-end differentiable network consists of two streams. The grid prediction stream takes as input reduced-resolution content \tilde{I}_c and style \tilde{I}_s images, learns the joint distribution between their low-level features, and predicts an affine bilateral grid Γ . The rendering stream, unmodified from HDRnet, operates at full-resolution. At each pixel (x, y, r, g, b), it uses a learned lookup table to compute a "luma" value z = g(r, g, b), slices out $A = \Gamma(x/w, y/h, z/d)$ (using trilinear interpolation), and outputs $O = A * (r, g, b, 1)^T$. By decoupling grid prediction resolution from that of rendering, our architecture offers a tradeoff between stylization quality and performance. Figure 3 summarizes the entire network and we describe each block below.

Grid Prediction. We aim to first learn a multi-scale model of the joint distribution between content and style features, and from this distribution, predict an affine bilateral grid. Rather than using strided convolutional layers to directly learn from pixel data, we follow recent work [10, 18, 8] and use a pretrained VGG-19 network to extract low-level features from both images at four scales ($conv1_1$, $conv2_1$, $conv3_1$, and $conv4_1$). We process these multi-resolution feature maps with a sequence of *feature transfer blocks* inspired by the StyleGAN architecture [11] (Figure 3). Starting from the finest level, each transfer block applies a stride-2 weight-sharing convolutional layer to both content and style features,

halving spatial resolution while doubling the number of channels (see Table 1). The shared-weight constraint crucially allows the following AdaIN layer to learn the joint content/style distribution without channel-wise correspondence supervision. Once the content feature map is normalized, we concatenate it to the similarly AdaIN-aligned feature maps from the pretrained VGG-19 layer of the same resolution. Since the content feature map now contains more channels, we use a stride-1 convolutional layer to select the relevant channels between learned-and-normalized vs. pretrained-and-normalized features.

We use three transfer blocks in our architecture, corresponding to the three finest-resolution layers of the selected VGG features. While using additional transfer blocks is possible, they are too coarse and replacing them with standard stride-2 convolutions makes little difference in our experiments. Since this component of the network effectively learns the relevant bilateral-space content features based on its corresponding style, it can be thought of as *learned feature transfer*.

Joint Bilateral Learning. With aligned-to-style content features in bilateral space, we seek to learn an affine bilateral grid that encodes a transformation that locally captures style and is aware of scene semantics. Like HDRnet, we split the network into two asymmetric paths: a fully-convolutional *local path* that learns local color transforms and thereby sets the grid resolution, and a *global path*, consisting of both convolutional and fully-connected layers, that learns a summary of the scene and helps spatially regularize the output transforms. The local path consists of two stride 1 convolutional layers, keeping the spatial resolution and number of features constant. This provides enough depth to learn local affine transforms without letting its receptive field grow too large (and thereby discarding any notion of spatial position).

As we aim to perform universal style transfer without any explicit notion of semantics (e.g., per-pixel masks provided by an external pretrained network), we use a small network to learn a global notion of scene category. Our global path consists of two stride 2 convolutional layers to further reduce resolution, followed by four fully-connected layers to produce a 64-element vector "summary". We concatenate the summary at each x, y spatial location output from the local path and use a 1×1 convolutional layer to reduce the final output to 96 channels. These 96 channels can be reshaped into a 8 "luma bins" that separate edges, each storing a 3×4 affine transform. We use the ReLU activation after all but the final 1×1 fusion layer and zero-padding for all convolutional layers.

2.3 Losses

Since our architecture is fully differentiable, we can simply define our loss function on the generated output. We augment the content and style fidelity losses of Huang et al. [8] with a novel *bilateral-space Laplacian regularizer*, similar to the one in [6]:

$$\mathcal{L} = \lambda_c \mathcal{L}_c + \lambda_{sa} \mathcal{L}_{sa} + \lambda_r \mathcal{L}_r, \tag{5}$$

where \mathcal{L}_c and \mathcal{L}_{sa} are the content and style losses defined in Equations 2 and 4, and

$$\mathcal{L}_r(\Gamma) = \sum_s \sum_{t \in N(s)} ||\Gamma[s] - \Gamma[t]||_F^2, \tag{6}$$

where $\Gamma[s]$ is one cell of the predicted bilateral grid, and $\Gamma[t]$ one of its neighbors.

The Laplacian regularizer penalizes differences between adjacent cells of the bilateral grid (indexed by s and finite differences computed over its six-connected neighbors N(s)) and encourages the learned affine transforms to be smooth in both space and intensity [2, 6]. As we show in our ablation study (Sec 3.1), the Laplacian regularizer is necessary to prevent visible grid artifacts.

We set $\lambda_c = 0.5$, $\lambda_{sa} = 1$, and $\lambda_r = 0.15$ in all experiments.

2.4 Training

We trained our model on high-quality landscape photos using Tensorflow [1], without any explicit notion of semantics. We use the Adam optimizer [12] with hyperparameters $\alpha = 10^{-4}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, and a batch size of 12 content/style pairs. For each epoch, we randomly split the data into 50000 content/style pairs. The training resolution is 256×256 and we train for a fixed 25 epochs, taking two days on a single NVIDIA Tesla V100 GPU with 16 GB RAM. Even trained at a low-resolution, our algorithms still performs well with 12 megapixel inputs, as shown in Figure 8. We attribute this to the fact that our losses are derived from pretrained VGG features, which are relatively invariant with respect to resolution.

3 Results

For evaluation, we collect a test set of 400 high-quality photos of a variety of subjects. We compare our algorithm to the state of the art in photorealistic style transfer, and conduct a user study. Furthermore, we perform a set of ablation studies to better understand the contribution of various components. Detailed comparisons with high-resolution images are included in the supplement.

3.1 Ablation Studies

Feature Transfer. We conduct multiple ablations to show the importance of our feature transfer blocks S.

First, we consider replacing S with two baseline networks: AdaIN [8] or WCT [17]. Starting with the same features extracted from VGG-19, we perform feature matching using AdaIN or WCT. The rest of the network is unchanged: that is, we attempt to learn local and global features directly from the baseline encoders and predict affine bilateral grids. Figure 4 (b) is the result of using the pretrained AdaIN feature from the "top path" only. The results in Figure 4 (b) and (c) show that while content is preserved, there is an overall color cast as well

9



Fig. 4: Ablation studies on feature transfer blocks. (a)-(e): We demonstrate the importance of our feature transfer architecture by replacing it with baseline networks. (f)-(j): Visualization of the contribution of each transfer block by disabling statistical feature matching on the others.

as inconsistent blotches when using either AdaIN or WCT. The low-resolution features simply lack the information density to learn even global color correction. Adding our multi-scale feature transfer blocks resolves this issue.

Second, to illustrate the contribution of each transfer block, we visualize our network's output when all but one block is disabled (including the top path inputs). As shown in Figure 4(f-j), earlier, finer resolution blocks learn texture and local contrast, while later blocks capture more global information such as the style input's dominant color tone, which is consistent with our intuition. By combining all transfer blocks at three different resolutions, our model merges these features at multiple scales into a joint distribution.

Network component ablations. To demonstrate the importance of other network components, in Figure 5, we compare our architecture with three variants: one trained without the bilateral-space Laplacian regularization loss (Equation 4), one without the global scene summary (Figure 3, yellow block), and one without "top path" inputs (Figure 3, dark green block). We also show that our network learns stylization parameterized as local affine transforms.

Figure 5 (b) shows distinctive dark halos when bilateral-space Laplacian regularization is absent. This is due to the fact that the network can learn to set



(a) Inputs (b) No \mathcal{L}_r (c) No summary (d) No top path (e) Full results

Fig. 5: Network component ablations.

regions of the bilateral grid to zero where it does not encounter image data (because images occupy a sparse 2D manifold in the grid's 3D domain). When sliced, the result is a smooth transition between zero and the proper transform.

Figure 5 (c) shows that the global summary helps with spatial consistency. For example, in the *mountain* photo, the left part of sky is saturated while the right side of the mountain is slightly washed out. In contrast, the output of our full network in Figure 5 (e) has more spatially consistent color. This is consistent with the observation in Gharbi et al. [5].

Figure 5 (d) demonstrates the necessity and importance of pretrained-andnormalized features (the output of the "top path" in Figure 3). The results show distinctive patches of incorrect color characteristic of the network locally overfitting to the style input. Adaptively selecting between learned and pretrained features at multiple resolutions eliminates this inconsistency.

Finally, we show that our network is not just an edge-aware interpolator—it learns stylization parameterized as local affine transforms. We set an edge-aware interpolation baseline by first running the full AdaIN network [8] on our 256×256 content and style images to produce a low-resolution stylized result. We then use BGU [2] to fit a $16 \times 16 \times 8$ affine bilateral grid (the same resolution as our network), and slice it with the full-resolution input to produce a full-resolution output. Figure 4 (d) shows that this strategy works quite poorly: AdaIN's output exhibits spatial distortions even at 256×256 , which BGU cannot fix.

Grid Spatial Resolution. Figure 6 (top) shows how the spatial resolution of the grid affects stylization quality. By fixing the number of luma bins at 8, the 1×1 case is a single global curve, where the network learns an incorrectly colored compromise. Going up to 2×2 , the network attempts to spatially vary the transformation, with slightly different colors applied to each quadrant, but the result is still an unsatisfying tradeoff. At 8×8 , there is sufficient spatial resolution to produce a faithful stylization result.

Grid Luma Resolution. Figure 6 (bottom) also shows how the "luma" resolution affects stylization quality, with a fixed spatial resolution 16×16 . With 1



Style 16x16x1 16x16x2 16x16x8



Fig. 6: Output using grids with different spatial (top) or luma (bottom) resolutions ($w \times h \times luma$ bins).

Fig. 7: Our method is robust to adversarial inputs such as when the content image is a portrait (an unseen category) or even a monochromatic "style".

Inputs	Output	Zoomed-in detail

(a) Output at 12 megapixels.

mage Size	I HOLOW CI	LOI	WOI	Ours
512×512	0.68s	0.25s	3.85s	< 5 ms
1024×1024	1.51s	0.84s	6.13s	< 5 ms
1000×2000	2.75s	OOM	10.94s	< 5 ms
2000×2000	OOM	OOM	OOM	< 5 ms
3000×4000	OOM	OOM	OOM	< 5 ms

Income Sine | Dhoto WOT | I ST | WOT² | Ound

(b) NVIDIA Tesla V100 runtime.

Mean Score	PhotoWCT	LST	WCT^2	Ours
Photorealism	2.02	2.89	4.21	4.14
Stylization	3.10	3.19	3.24	3.49
Overall quality	2.23	2.84	3.60	3.79

(c) User study results (higher is better).

Fig. 8: (a) The output of our method running at 12 megapixels, a typical smartphone camera resolution. Despite being trained at a fixed low resolution, our method produces sharp results while faithfully transferring the style from a significantly different scene. See the supplement for full-resolution images. (b) Performance benchmarks on a NVIDIA Tesla V100 GPU with 16 GB of RAM. OOM indicates out of memory. Note that photorealistic postprocessing adds significant overhead to LST performance. Due to GPU startup and memory I/O time, we were unable to get a precise measurement below 5ms. (c) Mean user study scores from 1200 responses. Raters scored the three output images in each sextet on a scale of 1-5 (higher is better).

luma bin, the network is restricted to predicting a single affine transform per tile. Interpolating between 2 luma bins reduces to a quadratic spline per tile, which is still insufficient for this image. In our experiments, 8 luma bins is sufficient for most images in our test set.

Qualitative Results 3.2

Visual Comparison. We compare our technique against three state-of-the-art photorealistic style transfer algorithms: PhotoWCT [18], LST [14], and WCT² [32],

using default settings. Note that for PhotoWCT, we use the most recent version of NVIDIA's FastPhotoStyle code. Comparisons with other algorithms are included in the supplementary material.

Figure 10 features a small sampling of the test set with some challenging examples. Owing to its reliance on unpooling and postprocessing, PhotoWCT results contain noticeable artifacts on nearly all scenes. LST mainly focuses on artistic style transfer, and to generate photorealistic results, uses a computeintensive spatial propagation network as a postprocessing step to reduce distortion artifacts. Figure 10 shows that there are still noticeable distortions in several instances, even after postprocessing. WCT² performs quite well when content and style are semantically similar, but when the scene content is significantly different from the landscapes on which it was trained, the results appear "hazy". Our method performs well even on these challenging cases. Thanks to its restricted output space, our method always produce sharp images which degrades gracefully towards the input (e.g., face, leaves) when given inputs outside the training set. Our primary artifact is a noticeable reduction in contrast along strong edges and is a known limitation of the local affine transform model [2].

Robustness. Thanks to its restricted transform model, our method is significantly more robust than the baselines when confronted with adversarial inputs, as shown in Figure 7. Although our model was trained exclusively on landscapes, the restricted transform model allows it to degrade gracefully on portraits which it has never encountered and even a monochromatic "style".

3.3 Quantitative Results

Runtime and Resolution. As shown in Figure 8 (b), our runtime on a workstation GPU significantly outperforms the baselines and is essentially invariant to resolution at practical resolutions. This is due to the fact that grid prediction, the "deep" part of the network, runs at a constant low resolution of 256×256 . In contrast, our full-resolution stream does minimal work and has hardware acceleration for trilinear interpolation. On a modern smartphone, inference runs comfortably above 30 Hz at full 12 megapixel camera resolution when quantized to 16-bit floating point and work distributed between ML acceleration hardware and the GPU. Figure 8 shows one such example. More images and a detailed performance benchmark are included in the supplement.

User Study. The question of whether an image is a faithful rendition of the style of another is inherently a matter of subjective taste. As such, we conducted a user study to judge whether our method delivers subjectively better results compared to the baselines. We recruited 20 users unconnected with the project. Each user was shown 20 sextets of images consisting of the input content, reference style, and four randomly shuffled outputs (PhotoWCT [18], WCT² [32], LST [14], and ours). For each output, they were asked to rate the following questions on a scale of 1–5:

Joint Bilateral Learning for Real-time Universal Photorealistic Style Transfer

- How noticeable are artifacts (i.e., less photorealistic) in the image?
- How similar is the output in style to the reference?
- How would you rate the overall quality of the generated image?

In total, we collected 1200 responses (400 images \times 3 questions) and results are shown in Figure 8 (c). In both stylization and overall quality, our technique outperforms the baselines: PhotoWCT, LST, and WCT². In terms of photorealism, our algorithm is on-par with WCT², and is significantly better than PhotoWCT and LST.

Video Stylization. Although our network is trained exclusively on images, it generalizes well to video content. Figure 9 shows an example where we transfer the style of a single photo to a video sequence that varies dramatically in appearance. The resulting video has a consistent style and is temporally coherent without any additional regularization or data augmentation.



Fig. 9: Transferring the style of a still photo to a video sequence. Although the content frames undergo substantial changes in appearance, our method produces a temporally coherent result consistent with the reference style. Please refer to the supplementary material for the full videos.

4 Conclusion

We presented a feed-forward neural network for universal photorealistic style transfer. The key to our approach is using deep learning to predict affine bilateral grids, which are compact image-to-image transformations that implicitly enforce the photorealism constraint. We showed that our technique is significantly faster than state of the art, runs in real-time on a smartphone at high resolution, and degrades gracefully even in extreme cases. We believe its robustness and fast runtime will lead to practical applications in mobile photography. As future work, we hope to further improve performance by reducing network size, and investigate how to relax the photorealism constraint to generate a continuum between photorealistic images and abstract art.

13



Inputs PhotoWCT [18] LST [14] WCT² [32] Ours Fig. 10: **Qualitative comparison** of our method against three state of the art baselines on some challenging examples.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al.: Tensorflow: A system for large-scale machine learning. In: 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI) (2016)
- Chen, J., Adams, A., Wadhwa, N., Hasinoff, S.W.: Bilateral guided upsampling. ACM TOG (2016)
- Dumoulin, V., Shlens, J., Kudlur, M.: A learned representation for artistic style. ICLR (2017)
- Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: CVPR (2016)
- 5. Gharbi, M., Chen, J., Barron, J.T., Hasinoff, S.W., Durand, F.: Deep bilateral learning for real-time image enhancement. ACM TOG (2017)
- Gupta, M., Cotter, A., Pfeifer, J., Voevodski, K., Canini, K., Mangylov, A., Moczydlowski, W., van Esbroeck, A.: Monotonic calibrated interpolated look-up tables. Journal of Machine Learning Research (2016)
- He, M., Liao, J., Chen, D., Yuan, L., Sander, P.V.: Progressive color transfer with dense semantic correspondences. ACM TOG (2019)
- 8. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV (2017)
- Jing, Y., Yang, Y., Feng, Z., Ye, J., Yu, Y., Song, M.: Neural style transfer: A review. TVCG (2019)
- 10. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: ECCV (2016)
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: CVPR (2019)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
- 13. Laffont, P.Y., Ren, Z., Tao, X., Qian, C., Hays, J.: Transient attributes for highlevel understanding and editing of outdoor scenes. ACM TOG (2014)
- 14. Li, X., Liu, S., Kautz, J., Yang, M.H.: Learning linear transformations for fast image and video style transfer. In: CVPR (2019)
- Li, Y., Wang, N., Liu, J., Hou, X.: Demystifying neural style transfer. In: IJCAI (2017)
- 16. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Diversified texture synthesis with feed-forward networks. In: CVPR (2017)
- 17. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Universal style transfer via feature transforms. In: NeurIPS (2017)
- Li, Y., Liu, M.Y., Li, X., Yang, M.H., Kautz, J.: A closed-form solution to photorealistic image stylization. In: ECCV (2018)
- Liu, S., De Mello, S., Gu, J., Zhong, G., Yang, M.H., Kautz, J.: Learning affinity via spatial propagation networks. In: NeurIPS (2017)
- Luan, F., Paris, S., Shechtman, E., Bala, K.: Deep photo style transfer. In: CVPR (2017)
- 21. Paris, S., Durand, F.: A fast approximation of the bilateral filter using a signal processing approach. In: ECCV (2006)
- 22. Pitié, F., Kokaram, A.C., Dahyot, R.: N-dimensional probability density function transfer and its application to color transfer. In: ICCV (2005)

- 16 X. Xia et al.
- 23. Puy, G., Pérez, P.: A flexible convolutional solver for fast style transfers. In: CVPR (2019)
- 24. Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. IEEE Computer Graphics and Applications (2001)
- 25. Shih, Y., Paris, S., Barnes, C., Freeman, W.T., Durand, F.: Style transfer for headshot portraits. ACM TOG (2014)
- 26. Shih, Y., Paris, S., Durand, F., Freeman, W.T.: Data-driven hallucination of different times of day from a single outdoor photo. ACM TOG (2013)
- 27. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognitio. In: ICLR (2015)
- 28. Tsai, Y.H., Shen, X., Lin, Z., Sunkavalli, K., Yang, M.H.: Sky is not the limit: semantic-aware sky replacement. ACM TOG (2016)
- Ulyanov, D., Lebedev, V., Vedaldi, A., Lempitsky, V.S.: Texture networks: Feedforward synthesis of textures and stylized images. In: ICML (2016)
- Ulyanov, D., Vedaldi, A., Lempitsky, V.: Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In: CVPR (2017)
- Wu, F., Dong, W., Kong, Y., Mei, X., Paul, J.C., Zhang, X.: Content-Based Colour Transfer. Computer Graphics Forum (2013)
- 32. Yoo, J., Uh, Y., Chun, S., Kang, B., Ha, J.W.: Photorealistic style transfer via wavelet transforms. In: ICCV (2019)