




Generative Adversarial Shaders for Real-Time Realism Enhancement

A. Salmi , Sz. Cséfalvay  and J. Imber 

Imagination Technologies, United Kingdom
{Arturo.Salmi, Szabolcs.Csefalvay, James.Imber}@imgtec.com



Figure 1: Our proposed Generative Adversarial Shaders combine the quality and automation of an adversarial training methodology with the computational efficiency and performance of a real-time post-processing shader pipeline. This crop is from a 1080p frame which was processed in 2.8ms on an Imagination Technologies BXT-32-1024 MC4 GPU to match the style of the Cityscapes dataset [COR* 16].

Abstract

Application of realism enhancement methods, particularly in real-time and resource-constrained settings, has been frustrated by the expense of existing methods. These achieve high quality results only at the cost of long runtimes and high bandwidth, memory, and power requirements. We present an efficient alternative: a high-performance, generative shader-based approach that adapts machine learning techniques to real-time applications, even in resource-constrained settings such as embedded and mobile GPUs. The proposed learnable shader pipeline comprises differentiable functions that can be trained in an end-to-end manner using an adversarial objective, allowing for faithful reproduction of the appearance of a target image set without manual tuning. The shader pipeline is optimized for highly efficient execution on the target device, providing temporally stable, faster-than-real time results with quality competitive with many neural network-based methods.

CCS Concepts

• **Computing methodologies** → **Computer graphics; Machine learning; Image processing; Rendering; Graphics processors;**

1. Introduction

Achieving greater photorealism has been a core concern of computer graphics research almost since the inception of the field. However, despite significant advances due to factors such as improved hardware capabilities, rendering algorithms, material representations and artistic tooling, there remains a stubborn shortfall in photorealism in real-time applications. This is particularly pronounced in settings with limited resources (e.g. bandwidth, compute, and power), such as smartphone GPUs.

There is a long history of using post-processing techniques to enhance the perceived realism of renders, for example by using programmable shaders to add such effects as chromatic aberration, lens flare, film grain, bokeh and motion blur to rendered images [MRD12] [GMN14]. It is notable that many such techniques mimic artifacts of the image capture process rather than natural appearance. These methods are designed and tuned by hand in what can be a labor-intensive process, and the final result is often as much a question of artistic preference as mimicking a physical camera.

A compelling alternative approach to post-processed realism enhancement is afforded by conditional generative models. Richter et al. [RAK21] demonstrate that image-to-image translation techniques utilizing neural networks and adversarial training [Izze17] can bridge the realism gap without the introduction of artifacts present in comparable methods [PLWZ19]. Furthermore, the appearance of the final renders can be matched to that of a target dataset. Their method achieves an impressively high degree of realism and visual quality, but would be prohibitively computationally expensive for real-time deployment on datacenter, desktop and embedded systems, particularly those without AI accelerators.

We hypothesize that much of the improvement in realism achieved by such deep learning-based methods may be accounted for by means of a carefully engineered shader pipeline exploiting domain knowledge, in particular by modelling the image capture process. A key advantage of conditional GANs [PEZZ20] [JZH*20] [RAK21] over such a solution is that the manual tuning process is replaced with numerical optimization. This suggests a hybrid approach that combines the computational efficiency of a bespoke post-processing shader pipeline with the automation and quality of an adversarial training methodology. We also take inspiration from the physical capture models [HVK*21] that are occasionally used in the tuning of image signal processors.

With the objective of preserving the learning capabilities of neural networks while designing a model that can be used for real-time rendering, we therefore propose a method based on a learnable post-processing pipeline composed of lightweight functions with a limited parameter set that we term Generative Adversarial Shaders (GAS) (Figure 1). By designing these functions to be differentiable, the pipeline can be trained in the same way as a neural network. We show that our method can be trained in an adversarial framework to learn a mapping between renders and photographic images. This provides several advantages over more conventional deep generative models:

- Improved computational efficiency over an undifferentiated neural network. Application of domain knowledge reduces the number of parameters by several orders of magnitude. The number

of operations is brought well within the range that can be processed in real time by an embedded GPU alongside a rendering workload.

- High degree of interpretability, debuggability and explainability due to its modular structure and low parameter count. Each stage in the pipeline has a clearly defined task by virtue of the function it performs. Parameters are typically semantically meaningful.
- Reduced scope for overfitting and introduction of artifacts by virtue of the limited functional range of the solution. The flip side of this is that we are not able to model more complex effects such as volumetric foliage. We suggest that these can be added by inclusion of a small auxiliary neural network operating alongside the proposed method, but this is left to future work.

The resulting shader pipeline (Figure 2) can then, for example, be deployed in a conventional manner, for example in OpenCL on a GPU. We only require final renders as inputs, which improves the ease of integration of our method. We do not require metadata such as G-buffer renders. The results we obtain are a step in the direction of real-time photorealism on resource-constrained platforms.

2. Related Work

Physically-based methods are a widespread approach for synthesizing photorealistic images. These simulate the physical processes that are involved in image formation, such as light transport [DBB06], [MTAS01] and other physical effects such as translucent media [HL01]. The quality of the images synthesized by these methods depends on the approximation of the physical models used: accurate modelling will produce realistic results, but these also typically come with large costs in terms of computational requirements and technical and artistic labor. Approximations of these methods can obtain real-time or interactive-time methods with an acceptable level of realism [REKS13].

Conditional image synthesis methods (for example, conditional GANs) can produce photorealistic images given a corresponding input image in a different domain, commonly referred to as *image-to-image translation* [Izze17] [PEZZ20] [JZH*20] [HLBK18].

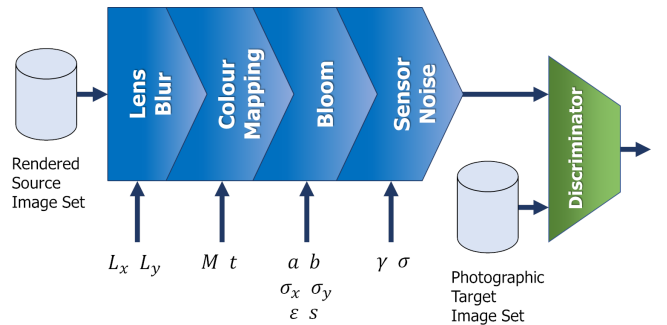


Figure 2: In our approach, the proposed differentiable post-processing pipeline (blue) comprising Generative Adversarial Shaders is trained as a generator in a conditional GAN. After training, the pipeline with learned parameters is deployed in a performance-optimized form (in our case, as OpenCL kernels).

Recent advances in diffusion models and attention-based architectures for conditional image generation have also been applied to the image-to-image translation problem [THY*23] [PSZ*23]. Such image-to-image methods are applicable to post-processed realism enhancement of rendered images [HTP*18]. Park et al. [PLWZ19] extend the idea to reliable generation of realistic synthetic images conditioned on semantic label maps; however when applied to the realism enhancement problem this leaves the problem largely unconstrained, reducing artistic control and in some cases impacting quality.

Despite giving plausible results in many cases, deep image-to-image models are prone to jarring artifacts such as hallucinations which would be intolerable for many applications (noted by Richter et al. [RAK21]). By using metadata extracted from the G-buffers to provide more scene information to the generator and further constrain the image formation process, Richter et al. overcome these recurrent issues and achieve impressive results both in terms of spatial and temporal stability and level of realism achieved. However, the proposed neural network is too expensive in terms of computational requirements and execution time to be feasible for real-time rendering applications.

Other techniques apply a data-driven approach to style transfer and realism enhancement. Early work showed that matching color distributions of renders to photographic examples can enhance the realism of computer graphics images [RAGS01], [PKD05], [PKD07]. With the adoption of CNNs, more sophisticated techniques for style transfer have been proposed [GEB16], [HB17], [ZAQW20]. Recent work has also focused on adopting an encoder/decoder structure to transfer the style of a target image while maintaining the content of the source image [LLL*18], [YUC*19]. Li et al. [LLKY19] propose a linear transformation method, where the mapping is locally modulated by a controller CNN. These approaches often depend upon a favorable overlap between the source and the target images. Large differences in content or layout can be detrimental for the quality of the results produced by these methods. This makes these methods less suitable for dynamic scene rendering and application to a large variety of rendered scenes.

Our proposed method combines traditional post-processing techniques with an adversarial training approach. In common with Carlson et al. [CSVJR19] we identify modelling of the image capture process with an appropriate series of post-processing stages as being a key component of realism enhancement. However, we go further in proposing that our pipeline be differentiable, with the parameters directly optimized to reduce an adversarially defined loss. Nalbach et al. [NAM*17] proposed a learnable post-processing method where the entire post-processing pipeline is approximated using a "Deep Shader" neural network. Instead, we design each shader as a bespoke learnable function, so that the pipeline can be trained in the same way as a neural network while retaining explainability and computing performance. Work from Thomas et al. [TF17] approximates global illumination using conditional generative adversarial networks with the goal of achieving offline rendering quality at interactive rates. While both adopt adversarial training, our method differs from Thomas et al.'s approach as it is designed to replicate post-processing effects applied in screen-space.

3. Method

Our proposed method has a learnable post-processing pipeline at its core (Figure 2). Each shader in the pipeline is designed to replicate an aspect of the physical image capture process. The behavior of each shader is governed by a corresponding set of parameters which, due to the differentiability of the pipeline, can be optimized during training in a conventional manner using backpropagation of errors. We replace the generator in a conventional conditional GAN with this pipeline. In this way, each shader is trained to optimally replicate the relevant aspect of the target dataset.

The shader pipeline is implemented in a deep learning framework such as PyTorch. However, we note that while this implementation is convenient for training, this is not optimal for performance, particularly on a resource-constrained target platform. For this reason, we also implement the shaders as functionally-equivalent OpenCL compute kernels optimized for inference speed. The stages in the pipeline can be merged at deployment to prevent excess system bandwidth consumption. After training, the learned parameters are applied to this deployed version of the shader pipeline. Our approach is robust against strong artifacts and hallucinations frequently observed in deep learning-based approaches, by virtue of our constraining the operation of each shader to a tightly-defined mathematical function. To further improve training stability, we also adopt regularization strategies which are described in more detail in section 3.2.

We fully expect the proposed training and deployment methodology to be extensible beyond the functionality modelled by the pipeline detailed below, as described in section 4.5.

The pipeline training code and OpenCL implementation are available at <https://github.com/james-imber/generative-adversarial-shaders>.

3.1. Learnable Shader Pipeline

Our shader pipeline consists of four stages, each taking as inputs learnable parameters which can be trained through backpropagation when implemented in a deep learning framework like PyTorch or TensorFlow. We fully expect the proposed pipeline to be extensible with more shaders implementing additional functionality (e.g. lens distortion, chromatic aberration, etc). The shaders are applied in approximately the order in which the corresponding physical process in image capture occurs.

3.1.1. Lens Blur

The optics of physical camera lenses introduce a certain amount of blur. This is modelled with a lens blur shader as the first module of the pipeline, which we implement as a learnable separable convolution (Equation 1), where both kernels are initialized as unit impulses to avoid introducing translation at the start of training. The window size is set to 5 since this was found to give a good balance between quality and performance in the deployed system.

$$I_{out} = (I_{in} * L_x) * L_y \quad (1)$$

In Equation 1, I_{in} and I_{out} are the input and output images, L_x and



Figure 3: Example of the learned lens blur applied to a rasterized frame from the GTA5 dataset [RVRK16]. The jagged edges of geometry are mitigated by the learned kernel. The foliage in the background also looks more similar to the foliage in the target dataset.

L_x and L_y are the horizontal and vertical blur kernels (represented with 5 values each), and $*$ is the convolution operation. During training, all weights of both kernels are updated. The weights of each kernel are normalized to sum to one before convolution. This is done to maintain output image intensity, as the shader should not be changing the overall brightness of the input frame. As well as replicating the blurring artifacts introduced by real world camera lenses, the learned convolutional operator also acts as a cheap anti-aliasing filter for point-sampled geometry. As we show in Figure 3, this shader can learn to convolve the image in order to mitigate jagged rendering artifacts at geometry edges in the input image.

3.1.2. Color Mapping

A post-processing pipeline for realism enhancement should transfer the color characteristics of photographic examples to the input frames. Rather than operating on a single source and target example [RAGS01], our color mapping shader learns a fixed transformation of the color distribution from the source to the target dataset. We learn a full affine mapping capable of modelling brightness and contrast adjustment, to avoid restricting to purely linear [HVK*21] [RAGS01] or translational [CSVJR19] mappings. Our proposed affine transformation (Equation 2), comprises a 3×3 matrix for the linear part of the mapping, and a 3-element vector for the translational part. All 12 parameters of the model are learnable.

$$\begin{bmatrix} R'_{x,y} \\ G'_{x,y} \\ B'_{x,y} \end{bmatrix} = \begin{bmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{bmatrix} \begin{bmatrix} R_{x,y} \\ G_{x,y} \\ B_{x,y} \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \\ t_3 \end{bmatrix} \quad (2)$$

Since this operates in a pixel-wise manner without reference to the local neighborhood, it has no impact on the general structure of the image. We initialize the transform with the identity function to aid convergence. Some examples of our learned color mappings are shown in Figure 4. Best results were obtained when the color mapping shader was trained in conjunction with the bloom shader.

3.1.3. Bloom

Our bloom shader approximates the behavior of cameras when capturing scenes with bright, saturated light sources. The bloom effect is commonly implemented as a post-processing stage to enhance perceived realism in graphics [JO04]. In captured images it can often be seen that saturated regions "bleed" over to neighboring pixels, causing a characteristic glowing effect. Our proposed bloom shader performs the following steps:

1. Isolate high intensity pixels (the glow sources) in the input image I_{in} by soft thresholding with a sigmoid (Equation 3) to obtain a "glow mask" I_{glow} .

$$I_{glow} = I_{in} \cdot \left(\frac{1}{1 + e^{-b(I_{luma} - a)}} \right) \quad (3)$$

In common rendering pipelines, I_{glow} can be obtained by setting a minimum HDR value threshold, with pixels lower than this threshold set to black. The render dataset we used for training, however, only provides 8-bit tone-mapped SDR frames. Furthermore, we require that the threshold be learnable. To solve this issue, we use a sigmoid as the thresholding function applied to a luma channel I_{luma} from the input frame I_{in} . a and b are learnable parameters.

2. Apply a parameterized blur using a separable convolution kernel then blend with the input image (Equation 4) to obtain I_{bloom} .

$$I_{bloom} = I_{in} + \left(I_{glow} * G_x(\sigma_x^2) \right) * G_y(\sigma_y^2) \quad (4)$$

$G_x(\sigma_x^2)$ and $G_y(\sigma_y^2)$ are discrete 1D Gaussian kernels normalized to sum to one. The variances σ_x^2 and σ_y^2 of the kernels are learned separately. By doing so, the shader can learn different amounts of "glow" to apply horizontally and vertically.

3. Apply a non-linear colour mapping to I_{bloom} (Equation 5), obtaining I_{out} with the finished bloom effect applied.

$$I_{out} = \min \left(\left(\frac{e^{\epsilon s}}{e^{\epsilon s} - 1} \right) \left(1 - e^{-I_{bloom} \epsilon} \right), 1 \right) \quad (5)$$



Figure 4: The color mapping function transfers the global color characteristics of the target dataset to the rasterized input frames. As well as replicating the "bleeding" effect caused by overexposure, by separately processing bright regions, the bloom shader helps prevent them from distorting the learned color mapping.

Exposure ϵ is a learnable parameter that controls the steepness of the curve and s is a learnable saturation parameter that controls the saturation point of the curve.

Since bloom can have a large radius of influence, for efficiency we use a multi-resolution approach for generating the bloom map in steps 1 and 2 above, before upsampling to full resolution and summing to produce a single bloom map for blending in step 3. For this purpose, 4 bloom maps were generated: one each at full, half, quarter and one-eighth resolution. Best results are achieved by learning separate sets of parameters for each resolution of the bloom algorithm.

3.1.4. Sensor Noise

Camera sensors introduce noise during image capture. A commonly-used noise model [CSVJR19] [HVK*21] breaks this into two components: (1) pixel capture noise, which is a Poisson process P depending on intensity; and (2) normally distributed noise as a catch-all for other processes (such as read noise [Jan01] and thermal noise). The division by gain γ in the Poisson distribution in our model (Equation 6) is to correct for intensity scaling between input I_{in} and output I_{out} due to gain.

$$I_{out} = \gamma \left(P \left(\frac{I_{in}}{\gamma} \right) + \mathcal{N} \left(0, \sigma'^2 \right) \right) \quad (6)$$

For improved performance and ease of implementation, we approximate the Poisson-distributed noise with a zero-mean normal distribution. We also reparameterize the standard deviation of the

second term as $\sigma = \gamma \sigma'$ to remove the dependence on gain, which helps orthogonalise the model.

$$I_{out} = I_{in} + \mathcal{N} \left(0, \gamma I_{in} \right) + \mathcal{N} \left(0, \sigma^2 \right) \quad (7)$$

Correct modelling of the noise, particularly readout noise in dark regions, depends on the input image I_{in} being in linear color space. For this reason, we opt to apply inverse gamma before this shader, and reapply the gamma afterwards. The gain γ and standard deviation σ are the learnable parameters of this shader. To obtain valid gradients, we rearrange equation 7 to yield equation 8:

$$I_{out} = I_{in} + \sqrt{\gamma I_{in}} \mathcal{N} \left(0, 1 \right) + \sigma \mathcal{N} \left(0, 1 \right) \quad (8)$$

3.2. Training

The learnable shader pipeline (Section 3.1) acts as the generator in an adversarial learning framework. Since care has been taken to ensure that the pipeline is differentiable, the parameters of the pipeline are updated in a conventional manner using backpropagation of errors. Using the information provided by the adversarial loss, the pipeline parameters are optimized so that they minimize the difference between the post-processed and the real image sets. PatchGAN [LW16], [Izze17] was chosen as the discriminator, as by penalizing the pipeline on patches we can improve the local quality of effects such as bloom, as well as global details like color matching.

There is a gap in capacity and functional flexibility between the

Method		KID score (x1000)	FID score	Mean execution time
GTA [RVRK16]		76.88 ± 0.11	71.67	n/a
Color transfer	Color Transfer [RAGS01]	70.08 ± 0.15	66.12	0.183ms ± 0.042ms
	CDT [PKD07]	85.98 ± 0.08	76.91	1003ms ± 3.9ms
Photo style transfer	WCT2 [YUC*19]	62.71 ± 0.08	64.50	1650ms ± 5ms
	CyCADA [HTP*18]	27.98 ± 0.09	33.47	29.12ms ± 0.5ms
	CUT [PEZZ20]	36.68 ± 0.10	41.25	98.9ms ± 3ms
Image-to-Image Translation	MUNIT [HLBK18]	33.87 ± 0.07	41.87	167ms ± 5ms
	TSIT [JZH*20]	46.29 ± 0.10	48.80	227ms ± 2ms
	EPE † [RAK21]	31.57 ± 0.08	37.73	1268ms
Ours		59.67 ± 0.09	61.65	0.089ms ± 0.0013ms

Table 1: KID and FID scores measured for domain adaptation methods using Cityscapes [COR*16] as the target distribution. Execution times were measured on an Nvidia RTX 2080 GPU at 960x540 resolution, on which our proposed method (designed for 1080p on mobile GPUs) achieves well in excess of real-time performance. † These KID and FID scores were obtained on JPEG-compressed images, and execution time was estimated based on scaling in FLOPs from that reported in [RAK21].

shader pipeline and the discriminator. This can lead to training instability, as the discriminator has the potential to exploit inflexibility in the generator, making its task of distinguishing real from post-processed images easier. To limit this behavior, we added instance noise to our discriminator [SCT*16], [AB17], [MNG18], with the only difference that rather than adding noise to the logits of the discriminator we add noise between the layers. This has the effect of smoothing the distributions of the source and target datasets, making them more similar and hence harder to distinguish.

The images used during training are 256x256 random crops from the input data frames. This helps prevent the discriminator from memorizing the dataset. Before being fed into the discriminator, the edges of the pipeline output are cropped to prevent the discriminator from learning how to distinguish generated and target images based by any edge effects present in the target or post-processed image sets. Spectral normalization was used to stabilize the discriminator [MKKY18].

For the training loss, we adopted a relativistic discriminator loss inspired by the relativistic least-square GAN loss used by Jiang et al. [JGL*21] (following LSGAN [MLX*17]), as this was shown to perform well at matching local characteristics of the generated images in unpaired training settings.

The pipeline and the discriminator are trained using the ADAM optimizer with learning rate set to 10^{-4} for both models. The training is continued until convergence of the adversarial loss.

4. Results

We evaluate the performance of our model on the unsupervised domain adaptation task (specifically rasterized to photographic images), and compare it with other methods proposed to solve the same task. Specifically, we evaluate our method on mapping from the GTA5 dataset [RVRK16] to the Cityscapes dataset [COR*16].

4.1. Quantitative Evaluation

FID [HRU*17] and KID [BSAG18] are two commonly-used metrics used to evaluate the performance of generative models. Re-

cent work by Parmar et al. however showed that these scores can vary significantly depending on factors such as the type of pre-processing applied to the images used for evaluation and/or the compression scheme used on the image [PZZ22]. For reasons of reproducibility and stability, we used the evaluation methodology and code provided by Parmar et al., both for FID and KID, without modification. Furthermore, wherever possible the evaluation is conducted on images stored in lossless formats in order to avoid any score variation introduced by compression schemes.

Table 1 contains the scores and execution times of each method. The proposed method is faster than other methods in the evaluation, and in particular is orders of magnitude faster than deep-learning based methods (Figure 5). In addition to its fast execution time, it also achieves a respectable KID and FID score. We note, however, that KID and FID are forgiving of many unacceptable artifacts present in some of the deep-learning based methods, such as hallucinated scenery. Some deep learning models will introduce extraneous objects in the scene (such as trees in place of sky) which do not damage the measured score. These scores are also not indicative of temporal stability, where our algorithm excels.

4.2. Qualitative Evaluation

Outputs for the methods in this evaluation for the same input frame are given in Figure 9 for comparison. Compared to methods relying on good alignment between the input frame and the style target, our method is capable of generalizing to a large set of input frames without producing unusual artifacts or leaving the input frame unchanged due to bad overlapping between the frames. With the exception of EPE, other unpaired image-to-image methods have a marked tendency to hallucinate unwanted content such as floating trees, car logos and high frequency artifacts. While EPE achieves impressive results in terms of realism enhancement, on certain scenes it can produce a notable halo effect around edges (Figure 6), possibly related to the segmentation maps used as an input by the generator. To demonstrate the temporal stability of the trained pipeline, we provide a video to accompany this pa-

	Rasterized Input	Lens Blur	Color Mapping	Bloom	Noise
KID	76.88 ± 0.11	76.18 ± 0.12	67.31 ± 0.08	58.46 ± 0.09	59.67 ± 0.09
FID	71.67	71.55	65.39	59.05	61.65

Table 2: Ablation study of the proposed pipeline. Shaders are added one at a time from left to right, and the pipeline trained from scratch each time. The KID and FID score is measured on Cityscapes for the resulting enhanced images to analyze each shader’s contribution. The resulting scores are below the images.

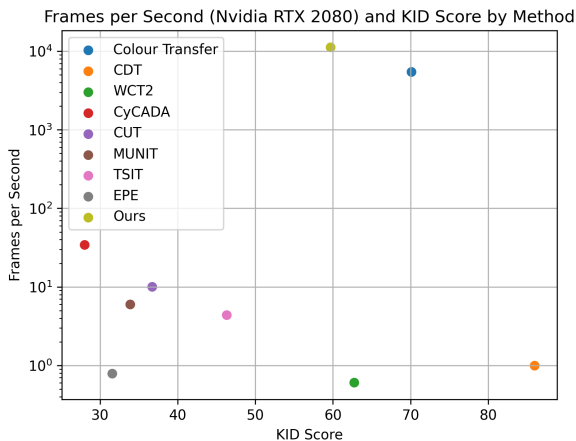


Figure 5: A graphical presentation of the data in Table 1. Note the logarithmic y-axis.

per [SCI23] including sequences from GTA 5 and from the *GamingVideoSet* dataset [BZS*18].

Without the need for auxiliary metadata such as scene G-buffers or semantic maps, our method manages to produce structurally coherent and stable scenes, both spatially and temporally, while also enhancing realism.

4.3. Ablation study

To evaluate the effectiveness of each shader in enhancing the realism of a rasterized input, we conducted ablation experiments in which the shaders were introduced to the pipeline one at a time, and the pipeline retrained from scratch each time. The KID and FID metrics were used to quantitatively evaluate the contribution of each shader. The results are shown in Table 2.

The addition of shaders to the pipeline progressively reduces both metrics, demonstrating that each shader improves the similarity between the enhanced images and the target dataset, with the exception of a small drop in score when the noise shader is introduced. We suspect that the reason for this is that when noise is introduced into the images, information is lost, affecting the features extracted by the Inception model used in both metrics.

4.4. Generalization to other Target Datasets and Input Content

To demonstrate the generality of our approach, our Generative Adversarial Shaders were trained on Mapillary Vistas [NOBK17] and KITTI [GLU12] in addition to Cityscapes. In Figure 7, we show some examples of the pipeline trained on KITTI and Mapillary Vistas. The trained pipeline can also generalize well to different types of input content without needing to be re-trained or fine-tuned on different input datasets. To demonstrate this, we applied the trained pipeline on different frames taken from the *GamingVideoSet* dataset [BZS*18]. The results are shown in Figure 8.



Figure 6: EPE produces halo artifacts around certain objects in the input frame. This is likely to be caused by the segmentation maps used by their proposed generator.



Figure 7: Output of the proposed pipeline trained on multiple photographic datasets with different characteristic styles, showing that our proposed method generalizes well. Insets show images sampled from the respective target datasets.

4.5. Limitations and Future Work

Our approach represents one end of what we expect is a continuum of methods ranging from real-time enhancement to high-quality, off-line realism enhancement. Prioritising computational performance

by using a series of handcrafted, learnable shaders instead of a large, undifferentiated neural network is a design choice that comes at the cost of limiting the range of effects our pipeline can replicate. For instance, unlike other models such as EPE [RAK21], our method cannot capture effects such as retexturing, introducing ad-



Figure 8: Output of the proposed trained pipeline when applied to different rasterised input content taken from GamingVideoSet [BZS*18], showing that our proposed method generalizes well to input content that has not been used during training.

ditional specularity or adding content in the input frame such as volumetric foliage. In future work these issues may be addressed by implementing "hybrid" techniques which combine the benefits of a bespoke shader pipeline with the learning capabilities of neural networks. To keep the pipeline lightweight and easily integrable, we avoided using additional input data such as depth maps and G-buffers. The use of this data may also be considered in future work as we believe it will also prove helpful when implementing such hybrid techniques.

Compared to other learning frameworks, adversarial training is known to be particularly unstable. In the case of our proposed method, such instabilities were exacerbated by the differences in learning capacities between the pipeline and the discriminator. In order to overcome these training limitations, adopting spectral normalization [MKKY18], instance noise [SCT*16], [AB17] and removing edge artifacts from the pipeline output by cropping is necessary when training the model. Furthermore, in order to avoid undesirable interactions which can further reduce training stability, the parameters of the pipeline are designed to be mutually orthogonal. This imposes a design constraint which is also often a feature of more conventional, hand-tuned image processing pipelines.

Our approach is focused on camera modelling as an important component for realism enhancement of rasterized frames. Although this may restrict the applications of the proposed pipeline to style-transfer tasks involving the realism enhancement of 3D content, we expect that similar bespoke pipelines can be constructed for other domains by designing shaders that exploit relevant domain understanding.

5. Conclusion

We have introduced Generative Adversarial Shaders, combining specialized, performance-optimized post-processing shader pipelines with adversarial training. Using a bespoke pipeline that draws on an understanding of image formation yields several advantages. The inefficient black box of an undifferentiated neural network is replaced with a computationally efficient, easily-understood, interpretable and modular white box. An orders-of-magnitude reduction in the number of learnable parameters reduces the functional range of the model, which makes it impervious to deep learning related artifacts such as hallucination, and aids with temporal stability. The improved efficiency enables real-time realism enhancement, even on resource-constrained devices running a graphics rendering workload. The fact that only SDR images are required for input also makes the proposed pipeline straightforward to integrate.

We have shown that camera modelling accounts for a significant part of the realism enhancement problem. In future this can be combined with neural networks to (1) improve the overall efficiency of a deep learning based solution and (2) account for screen-space geometric effects (such as volumetric foliage [RAK21]) that cannot easily be covered by such a shader pipeline. We also expect that additional shaders can be added to the pipeline to model other image capture effects and further improve realism.

Acknowledgements

We would like to thank our colleagues at Imagination Technologies for their help and support with this project. In particular, we would like to acknowledge Christian Burisch for sharing his thoughts on realism enhancement, and Kristof Beets for helpful discussions. We

thank the reviewers for providing us with constructive feedback which has strengthened the paper. Lastly, we would like to thank Dr. Nabajeet Barman from Kingston University for granting us access to *GamingVideoSet* at short notice.

References

- [AB17] ARJOVSKY M., BOTTOU L.: Towards principled methods for training generative adversarial networks. In *International Conference on Learning Representations* (2017). URL: https://openreview.net/forum?id=Hk4_qw5xe. 6, 9
- [BSAG18] BIŃKOWSKI M., SUTHERLAND D. J., ARBEL M., GRETTON A.: Demystifying MMD GANs. In *International Conference on Learning Representations* (2018). URL: <https://openreview.net/forum?id=r1lUOzWCW>. 6
- [BZS*18] BARMAN N., ZADTOOTAGHAI S., SCHMIDT S., MARTINI M. G., MÖLLER S.: Gamingvideoset: A dataset for gaming video streaming applications. In *Proceedings of the 16th Annual Workshop on Network and Systems Support for Games* (2018), NetGames '18, IEEE Press. 7, 9
- [COR*16] CORDTS M., OMRAN M., RAMOS S., REHFELD T., ENZWEILER M., BENENSON R., FRANKE U., ROTH S., SCHIELE B.: The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 3213–3223. 1, 6
- [CSVJR19] CARLSON A., SKINNER K. A., VASUDEVAN R., JOHNSON-ROBERSON M.: Sensor transfer: Learning optimal sensor effect image augmentation for sim-to-real domain adaptation. *IEEE Robotics and Automation Letters* 4, 3 (2019), 2431–2438. 3, 4, 5
- [DDB06] DUTRE P., BALA K., BEKAERT P.: *Advanced Global Illumination, 2nd Edition*. A K Peters, Natick, MA, 2006. 2
- [GEB16] GATYS L. A., ECKER A. S., BETHGE M.: Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016). 3
- [GLU12] GEIGER A., LENZ P., URTASUN R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)* (2012). 7
- [GMN14] GUERTIN J.-P., MCGUIRE M., NOWROUZSAHRAI D.: A Fast and Stable Feature-Aware Motion Blur Filter. In *Eurographics/ACM SIGGRAPH Symposium on High Performance Graphics* (2014), Wald I., Ragan-Kelley J., (Eds.), The Eurographics Association. doi: 10.2312/hpg.20141093. 2
- [HB17] HUANG X., BELONGIE S. J.: Arbitrary style transfer in real-time with adaptive instance normalization. *2017 IEEE International Conference on Computer Vision (ICCV)* (2017), 1510–1519. 3
- [HL01] HARRIS M. J., LASTRA A.: Real-time cloud rendering. In *Proc. Eurographics* (2001), pp. 76–84. 2
- [HLBK18] HUANG X., LIU M.-Y., BELONGIE S., KAUTZ J.: Multimodal unsupervised image-to-image translation. In *ECCV* (2018). 2, 6, 11
- [HRU*17] HEUSEL M., RAMSAUER H., UNTERTHINER T., NESSLER B., HOCHREITER S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems* (2017), Guyon I., Luxburg U. V., Bengio S., Wallach H., Fergus R., Vishwanathan S., Garnett R., (Eds.), vol. 30, Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/8a1d694707eb0fe65871369074926d-Paper.pdf. 6
- [HTP*18] HOFFMAN J., TZENG E., PARK T., ZHU J.-Y., ISOLA P., SAENKO K., EFROS A. A., DARRELL T.: Cycada: Cycle consistent adversarial domain adaptation. In *International Conference on Machine Learning (ICML)* (2018). 3, 6, 11
- [HVK*21] HANSEN P., VILKIN A., KRUSTALEV Y., IMBER J., TALAGALA D., HANWELL D., MATTINA M., WHATMOUGH P. N.: Isp4ml: The role of image signal processing in efficient deep learning vision systems. In *2020 25th International Conference on Pattern Recognition (ICPR)* (2021), pp. 2438–2445. doi:10.1109/ICPR48806.2021.9411985. 2, 4, 5
- [IZZE17] ISOLA P., ZHU J.-Y., ZHOU T., EFROS A. A.: Image-to-image translation with conditional adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017), pp. 5967–5976. doi:10.1109/CVPR.2017.632. 2, 5
- [Jan01] JANESICK J.: *Scientific Charge Coupled Devices*. SPIE Publications, 2001. 5
- [JGL*21] JIANG Y., GONG X., LIU D., CHENG Y., FANG C., SHEN X., YANG J., ZHOU P., WANG Z.: Enlightenment: Deep light enhancement without paired supervision. *IEEE Transactions on Image Processing* 30 (2021), 2340–2349. 6
- [JOO4] JAMES G., O'RORKE J.: *GPU Gems: Programming Techniques, Tips and Tricks for Real-Time Graphics*, 1st ed. Pearson Higher Education, 2004. 4
- [JZH*20] JIANG L., ZHANG C., HUANG M., LIU C., SHI J., LOY C. C.: TSIT: A simple and versatile framework for image-to-image translation. In *ECCV* (2020). 2, 6, 11
- [LLKY19] LI X., LIU S., KAUTZ J., YANG M.-H.: Learning linear transformations for fast arbitrary style transfer. In *IEEE Conference on Computer Vision and Pattern Recognition* (2019). 3
- [LLL*18] LI Y., LIU M.-Y., LI X., YANG M.-H., KAUTZ J.: A closed-form solution to photorealistic image stylization. In *Computer Vision – ECCV 2018* (2018), Ferrari V., Hebert M., Sminchisescu C., Weiss Y., (Eds.), Springer International Publishing, pp. 468–483. 3
- [LW16] LI C., WAND M.: Precomputed real-time texture synthesis with markovian generative adversarial networks. In *ECCV* (3) (2016), Leibe B., Matas J., Sebe N., Welling M., (Eds.), vol. 9907 of *Lecture Notes in Computer Science*, Springer, pp. 702–716. URL: <http://dblp.uni-trier.de/db/conf/eccv/eccv2016-3.html#LiW16>. 5
- [MKKY18] MIYATO T., KATAOKA T., KOYAMA M., YOSHIDA Y.: Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations* (2018). URL: <https://openreview.net/forum?id=B1QRgziT->. 6, 9
- [MLX*17] MAO X., LI Q., XIE H., LAU R. Y., WANG Z., SMOLLEY S. P.: Least squares generative adversarial networks. In *2017 IEEE International Conference on Computer Vision (ICCV)* (2017), pp. 2813–2821. doi:10.1109/TCCV.2017.304. 6
- [MNG18] MESCHEDER L., NOWOZIN S., GEIGER A.: Which training methods for gans do actually converge? In *International Conference on Machine Learning (ICML)* (2018). 6
- [MRD12] MCINTOSH L., RIECKE B., DIPOLA S.: Efficiently simulating the bokeh of polygonal apertures in a post-process depth of field shader. *Computer Graphics Forum* 31 (09 2012), 1810–1822. doi:10.1111/j.1467-8659.2012.02097.x. 2
- [MTAS01] MYSZKOWSKI K., TAWARA T., AKAMINE H., SEIDEL H.-P.: Perception-guided global illumination solution for animation rendering. In *Computer Graphics (SIGGRAPH-2001): Conference Proceedings* (2001), Fiume E., (Ed.), Annual Conference Series, ACM, pp. 221–230. URL: <http://www.mpi-sb.mpg.de/resources/aqm/dynenv/paper/sg2001-myszkowski.pdf>. 2
- [NAM*17] NALBACH O., ARABADZHIYSKA E., MEHTA D., SEIDEL H.-P., RITSCHEL T.: Deep shading: Convolutional neural networks for screen space shading. *Comput. Graph. Forum* 36, 4 (jul 2017), 65–78. URL: <https://doi.org/10.1111/cgfm.13225>, doi: 10.1111/cgfm.13225. 3
- [NOBK17] NEUHOLD G., OLLMANN T., BULÒ S. R., KONTSCHEIDER P.: The mapillary vistas dataset for semantic understanding of street scenes. In *2017 IEEE International Conference on Computer Vision*



Figure 9: Comparison between our pipeline and other methods. Colour Transfer, CDT and WCT2 fail to produce optimal results as they depend on favorable target images. TSIT, MUNIT, CUT and CyCADA hallucinate the content of the input frame by introducing undesired content such as floating trees and high frequency artifacts.

- (ICCV) (2017), pp. 5000–5009. doi:10.1109/ICCV.2017.534.7
- [PEZZ20] PARK T., EFROS A. A., ZHANG R., ZHU J.-Y.: Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision* (2020). 2, 6, 11
- [PKD05] PITIÉ F., KOKARAM A. C., DAHYOT R.: N-dimensional probability density function transfer and its application to color transfer. *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume I 2* (2005), 1434–1439 Vol. 2. 3
- [PKD07] PITIÉ F., KOKARAM A. C., DAHYOT R.: Automated colour grading using colour distribution transfer. *Computer Vision and Image Understanding* 107, 1 (2007), 123–137. Special issue on color image processing. URL: <https://www.sciencedirect.com/science/article/pii/S1077314206002189>, doi:<https://doi.org/10.1016/j.cviu.2006.11.011>. 3, 6, 11
- [PLWZ19] PARK T., LIU M.-Y., WANG T.-C., ZHU J.-Y.: Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019). 2, 3
- [PSZ*23] PARMAR G., SINGH K. K., ZHANG R., LI Y., LU J., ZHU J.-Y.: Zero-shot image-to-image translation. *arXiv preprint arXiv:2302.03027* (2023). 3
- [PZZ22] PARMAR G., ZHANG R., ZHU J.-Y.: On aliased resizing and surprising subtleties in gan evaluation. In *CVPR* (2022). 6
- [RAGS01] REINHARD E., ASHIKHMIN M., GOOCH B., SHIRLEY P.: Color transfer between images. *IEEE Computer Graphics and Applications* 21 (5) (sep 2001), 34–41. 3, 4, 6, 11
- [RAK21] RICHTER S. R., ALHAJJA H. A., KOLTUN V.: Enhancing photorealism enhancement. *arXiv:2105.04619* (2021). 2, 3, 6, 8, 9, 11
- [REKS13] REINHARD E., EFROS A., KAUTZ J., SEIDEL H.-P.: On visual realism of synthesized imagery. *Proceedings of the IEEE* 101 (09 2013), 1998–2007. doi:10.1109/JPROC.2013.2260711. 2
- [RVRK16] RICHTER S. R., VINEET V., ROTH S., KOLTUN V.: Playing for data: Ground truth from computer games. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14* (2016), Springer, pp. 102–118. 4, 6, 11
- [SCI23] SALMI A., CSÉFALVAY S., IMBER J.: Generative adversarial shaders for real-time realism enhancement, 2023. URL: <https://www.youtube.com/watch?v=vzVbO9upKjk>. 7
- [SCT*16] SØNDERBY C. K., CABALLERO J., THEIS L., SHI W., HUSZÁR F.: Amortised map inference for image super-resolution. *arXiv preprint arXiv:1610.04490* (2016). 6, 9
- [TF17] THOMAS M. M., FORBES A.: Deep illumination: Approximating dynamic global illumination with generative adversarial network. 3
- [THY*23] TORBUNOV D., HUANG Y., YU H., HUANG J., YOO S., LIN M., VIREN B., REN Y.: Uvcgan: Unet vision transformer cycle-consistent gan for unpaired image-to-image translation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (2023), pp. 702–712. 3
- [YUC*19] YOO J., UH Y., CHUN S., KANG B., HA J.-W.: Photorealistic style transfer via wavelet transforms. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2019). 3, 6, 11
- [ZAQW20] ZHU P., ABDAL R., QIN Y., WONKA P.: Sean: Image synthesis with semantic region-adaptive normalization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020), pp. 5103–5112. doi:10.1109/CVPR42600.2020.00515. 3