

Distilling Style from Image Pairs for Global Forward and Inverse Tone Mapping

Aamir Mustafa
University of Cambridge
UK
aamir.mustafa@cl.cam.ac.uk

Param Hanji
University of Cambridge
UK
param.hanji@cl.cam.ac.uk

Rafał K. Mantiuk
University of Cambridge
UK
rafal.mantiuk@cl.cam.ac.uk




Figure 1: Distribution of the latent color code representation for different frames. We see that frames with similar content and similar color retouching have been encoded close together in the latent space. All frames are taken from the BBC documentary *Blue Planet II* Episode “Green Seas”. Note that the model trained to encode the style of frames in 2 dimensions (see Sec. 3.4.2) is used in this plot for better visualization.

ABSTRACT

Many image enhancement or editing operations, such as forward and inverse tone mapping or color grading, do not have a unique solution, but instead a range of solutions, each representing a different style. Despite this, existing learning-based methods attempt to learn a unique mapping, disregarding this style. In this work, we show that information about the style can be distilled from collections of image pairs and encoded into a 2- or 3-dimensional vector. This gives us not only an efficient representation but also an interpretable latent space for editing the image style. We represent the global color mapping between a pair of images as a custom normalizing flow, conditioned on a polynomial basis of the pixel color. We show that such a network is more effective than PCA or VAE at encoding image style in low-dimensional space and lets us obtain an accuracy close to 40 dB, which is about 7-10 dB improvement over the state-of-the-art methods.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CVMP '22, December 1–2, 2022, London, United Kingdom

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9939-5/22/12.

<https://doi.org/10.1145/3565516.3565520>

KEYWORDS

tone mapping, inverse tone mapping, color mapping, normalizing flows, dimensionality reduction

ACM Reference Format:

Aamir Mustafa, Param Hanji, and Rafał K. Mantiuk. 2022. Distilling Style from Image Pairs for Global Forward and Inverse Tone Mapping. In *European Conference on Visual Media Production (CVMP '22)*, December 1–2, 2022, London, United Kingdom. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3565516.3565520>

1 INTRODUCTION

Several applications in image and video processing require mapping source to target colors, such as mapping High Dynamic Range (HDR) to Standard Dynamic Range (SDR) frames (tone mapping), SDR to HDR frames (inverse tone mapping), RAW to color-graded frames or retouched images. Tone mapping and inverse tone mapping have become particularly relevant with the introduction of HDR formats and standards, which necessitates mastering content separately for both SDR and HDR devices for best presentation.

The research on forward and inverse tone mapping dates several decades [Eilertsen et al. 2017b; Mantiuk et al. 2008; Reinhard et al. 2002; Stockham Jr 1972; Tumblin and Rushmeier 1993]. Most tone mapping methods propose hand-crafted recipes [Reinhard et al. 2002; Stockham Jr 1972; Tumblin and Rushmeier 1993], or optimization criteria [Mantiuk et al. 2008], which should produce the

most desirable image. Rather than following this line of research to find the right recipe, we learn the mapping from collections of (source, target) image pairs manually prepared by skilled color artists. With the advent of deep learning and large datasets, recent works treat the relationship between paired images as a one-to-one mapping, which they typically learn with Convolutional Neural Networks (CNNs). The notable conceptual flaw of this approach is that tone mapping, color grading, retouching, and inverse tone mapping do not have a single solution. If the task is given to several skilled color artists, each one is likely to produce a different result. This was shown both in early research on tone mapping [Yoshida et al. 2006], and in the Adobe-MIT 5k dataset [Bychkovsky et al. 2011], where each image was differently retouched by 5 photographers. Each color artist conveys a different style, and methods that automate this process should respect and preserve such intended style. Because the style information is not present in the input image, existing methods that learn a CNN based image-to-image mapping can learn at most a single style. The possibility of an indefinite number of equally plausible solutions, makes a CNN based image-to-image translation network incapable of modelling the intended style without an additional meta-data as style encoding.

To this end, in this work, we deviate from the conventional approach of learning an one-to-one mapping function from the source to the target domain. Instead, we propose an orthogonal direction of research to distill the artistic style from pairs of source and target images and encode it in a low dimensional *style* vector. More specifically, we present a customized conditional Invertible Neural Network (INN), which captures the one-to-many nature of the problem. Unlike other works, our lightweight architecture operates in a pixel-wise manner to map the target pixels to a latent style, given the source pixels as conditioning inputs. The proposed bi-directional training, lets us converge the pixel-wise mapping to a single low dimensional latent *style* vector as the meta data for the entire image. As shown in Fig. 1, our trained INN can map images of similar styles to neighboring points in the latent style space. The proposed method allows us to distill the underlying global mapping between the source and the target image pair into a latent representation vector at *inference* time.

The main benefit of extracting style is that it lets us efficiently encode very precise color mapping from a source to target image. Such an encoding could be used for simultaneous transmission of HDR and SDR content, where each HDR frame is accompanied by a 2–4 dimensional vector, which lets us reconstruct SDR frames with almost perfect accuracy (~40 dB). The encoding can also be used to easily edit the color mapping (color grading, retouching, tone-mapping) using just a pair of sliders instead of much more complex interfaces used for those tasks. Finally, if we wish to perform fully automated mapping in an “average” style, we can run inference on our network while setting the style vector to 0.

The main contributions of this work are:

- We show that it is necessary to distill a style from image pairs to effectively learn a highly accurate color mapping (~40 dB) and that the style can be encoded in a 2–4 dimensional vector.
- We propose an INN-based generative model that can learn the distribution of styles and encode it in a low-dimensional

vector. The model can be trained on moderately sized datasets (500 image pairs) and achieves a dramatic gain in performance of over 10 dB compared to the methods that disregard style.

- The utility of the color mapping with a style vector is demonstrated in video transmission, semi-automatic color-grading, and inverse tone mapping.

In this work, we distill the style between the source-target image pairs, which have been manually color graded by color artists. Such color mappings for video content are only global in nature (the same for all pixels in an image, $\mathcal{M} : \mathcal{R}^3 \rightarrow \mathcal{R}^3$) and do not contain any local (spatially varying) changes. However, our method can be easily extended to a manual local mapping, by splitting frames into a number of tiles, as done in [Eilertsen et al. 2015].

2 RELATED WORK

Our work addresses the problem of learning color mapping from pairs of images, found in tone-mapping [Bychkovsky et al. 2011; Gharbi et al. 2017; Rana et al. 2020], inverse tone mapping [Eilertsen et al. 2017a; Liu et al. 2020; Marnerides et al. 2018; Santos et al. 2020], image enhancement and automatic retouching [He et al. 2020; Kim et al. 2021; Park et al. 2018; Wang et al. 2019; Yan et al. 2016; Zeng et al. 2020]. The majority of existing methods attempt to learn one-to-one mapping from a large collection of image pairs, such as that found in MIT-Adobe FiveK dataset [Bychkovsky et al. 2011]. The mapping can be represented using polynomial basis functions [Yan et al. 2016], 3D LUTs with trainable weights [Zeng et al. 2020], a linear combination of representative colors [Kim et al. 2021], color-to-color multi-layer perceptron (MLP) [He et al. 2020], an encoder-decoder architecture [Chen et al. 2018; Ignatov et al. 2017; Wang et al. 2019; Yan et al. 2016], a bilateral grid of affine color transformation matrices [Gharbi et al. 2017], or a Reinforcement Learning (RL) policy that mimics the sequence of operations that a human expert would take [Park et al. 2018]. All those methods attempt to extract local and global features from an input image so that the mapping function can adapt to image content. We show that this information alone is insufficient to obtain a highly accurate mapping.

Several works address the problem of learning different styles of image mapping, but none attempt to achieve the same goals as our work. PieNet [Kim et al. 2020b] and StarEnhancer [Song et al. 2021] learn embedding of a style (or personalization) of image-to-image mapping. They assume that all images processed by a single expert (from the Adobe-MIT-5K dataset) share the same style. We found this assumption to be overly optimistic, as experts often vary style between images. For that reason, we assign a separate style vector to each image pair. This lets us obtain an accuracy close to 40 dB, while both methods report only about 25 dB on the same dataset. Neither of the methods attempts to reduce the dimensionality of the style vector to make it better suited to coding and manipulation of images.

3 METHODOLOGY

3.1 Problem formulation

We aim to model a global, spatially invariant color mapping between pairs of images or video frames. Furthermore, such a mapping should be controlled by a low-dimensional *style* vector. Let $\{\mathbf{x}^i, \mathbf{y}^i\}_{i=1}^N$ represent N images pairs. Each RGB pixel p in those pairs is related by a global mapping \mathcal{M} :

$$\mathbf{x}_p^i = \mathcal{M}(\mathbf{y}_p^i, \mathbf{z}^i), \quad (1)$$

where i is the image index and \mathbf{z}^i is the low-dimensional vector that encodes *style* parameters specific to each pair. Please note that we use \mathbf{y} for source and \mathbf{x} for target color to be consistent with the normalizing flow notation used later.

Such a mapping can be effectively expressed as a parametric function using Polynomial Color Correction (PCC) basis functions [Finlayson and Drew 1997; Finlayson et al. 2015]:

$$\mathbf{x}_p^i = C(\mathbf{y}_p^i) \cdot \mathbf{M}^i, \quad (2)$$

where $C(\mathbf{y}_p^i)$ is the polynomial basis function for the pixel \mathbf{y}_p^i :

$$C(\mathbf{y}_p^i) = C([r, g, b]) = [r, g, b, r^2, g^2, b^2, rg, gb, br, \dots, r^4, g^4, b^4], \quad (3)$$

and \mathbf{M}^i is the 34×3 style matrix (a representation of \mathbf{z}^i). We found that at least a 4th degree PCC with $34 \times 3 = 102$ parameters is needed for our diverse set of scenes (further discussed in Sec. 6.1). However, such a large latent representation is not interpretable and cannot be intuitively modified to edit the style. While the obvious approach is to use one of the dimensionality reduction techniques such as Principal Component Analysis (PCA) and Variational Autoencoder (VAE), we found them to be ineffective for our problem, as we will show in Sec. 4.3. Instead, we propose to use a custom adaptation of normalizing flows [Rezende and Mohamed 2015] to represent the style mapping.

3.2 Our approach

We cast the problem of finding a suitable approximation of \mathcal{M} from Eq. 1 as a generative modeling task using supervised learning. We learn a conditional mapping from a known latent distribution to the target domain using a parametric model g_θ with weights θ . After training, the learned network g_θ and extracted latent vector \mathbf{z}^i reproduce the required per-pixel mapping,

$$\mathbf{x}_p^i = g_\theta(\mathbf{z}^i; \mathbf{c}_p^i). \quad (4)$$

The conditioning vector \mathbf{c}_p^i could be equal to the source pixel value, $\mathbf{c}_p^i = \mathbf{y}_p^i$. However, we found that performance is much better if the conditioning vector contains the polynomial basis functions: $\mathbf{c}_p^i = C(\mathbf{y}_p^i)$. Intuitively, when the polynomial basis is supplied, the network only needs to learn the appropriate linear combination, similar to matrix \mathbf{M}^i from Eq. 2. This dimensionality expansion approach is akin to positional encoding, used to represent multi-dimensional functions in NeRF methods [Mildenhall et al. 2020; Tancik et al. 2020]. We also observed that adding global information to the conditioning vector, such as the image histogram or VGG features [Simonyan and Zisserman 2014] to \mathbf{c}_p^i has little effect, as shown in the ablation studies in Sec. 6.1.

3.3 The latent distribution model

To approximate the global mapping \mathcal{M} with a parametric per-pixel function. We choose the class of generative models called normalizing flows [Dinh et al. 2016; Kingma and Dhariwal 2018; Rezende and Mohamed 2015]. They employ invertible layers with tractable determinants of jacobian to learn a non-trivial mapping from a known latent to the target distribution. The bidirectionality of normalizing flows makes them particularly suitable since we wish to both (1) extract a latent *style* vector from a set of pixels (reverse discriminative pass) and (2) reconstruct a pixel from the given input and latent vector (forward generative pass). This enables us to manipulate the image style by converting it to a latent space as an intermediate step (see Fig. 8).

Given many samples from the target distribution X , the objective is to train g_θ using Maximum Likelihood Estimation (MLE) to convert them into samples of a latent distribution Z . If g_θ is invertible, the change of variables formula expresses the probability of a target sample $\mathbf{x} \sim X$ as a function of the probability of the transformed latent $\mathbf{z} \sim Z$:

$$p_X(\mathbf{x}|\mathbf{c}) = p_Z\left(g_\theta^{-1}(\mathbf{x}; \mathbf{c})\right) \cdot \left|\det \nabla g_\theta^{-1}(\mathbf{x}; \mathbf{c})\right|. \quad (5)$$

Here \mathbf{c} is a conditioning vector, which corresponds to the polynomial coefficients \mathbf{M}^i in our implementation (see Eq. 2).

If the latents follow a known distribution such as a standard normal or uniform, Eq. 5 gives us an exact expression for the probability of any data point. For efficient MLE training, g_θ should be differentiable, and the determinant of its jacobian should be easy to compute. In practice, a normalizing flow is a deep architecture consisting of invertible layers whose jacobians are diagonal, lower-triangular, or the identity matrix.

3.4 Network architecture

Our INN, depicted in Fig. 2, is composed of a series of 8 invertible blocks each consisting of (1) affine coupling layer, (2) random permutation and (3) batch normalization with ActNorm [Kingma and Dhariwal 2018]. The coupling layers [Dinh et al. 2016] increase the expressive power of the INN by incorporating complex, non-invertible sub-networks $s(\cdot)$ and $t(\cdot)$. To keep the total number of trainable parameters small, we use simple 2-layer MLPs for both $s(\cdot)$ and $t(\cdot)$.

3.4.1 Conditional coupling. Affine coupling [Dinh et al. 2016] is a key operation of most normalizing flow architectures that increases the expressibility of the network for unconditional generative modeling. This involves splitting the inputs \mathbf{u} of an intermediate layer into two equal parts \mathbf{u}_1 and \mathbf{u}_2 . The first, \mathbf{u}_1 is transmitted unchanged, while \mathbf{u}_2 goes through an affine transformation where parameters are some functions of \mathbf{u}_1 , realized through complex, non-linear subnetworks $s(\mathbf{u}_1)$ and $t(\mathbf{u}_1)$.

However, we are interested in modeling the conditional distribution $p_X(\mathbf{x}|\mathbf{c})$ as described in Eq. 5. For this, we use conditional coupling described in [Ardizzone et al. 2020; Lugmayr et al. 2020] by appending \mathbf{c} from Eq. 3 to the inputs of $s(\cdot)$ and $t(\cdot)$. Now, the output of the conditional coupling layer is given by the concatenated vector $[\mathbf{v}_1, \mathbf{v}_2]$ where




Figure 2: The network architecture used for our pixel-wise conditional INN. Here the sub-networks $s(\cdot)$ and $t(\cdot)$ are fully connected networks with 2 hidden layers each. We use 8 invertible blocks in our architecture.

$$\mathbf{v}_1 = \mathbf{u}_1, \quad \mathbf{v}_2 = \mathbf{u}_2 \odot \exp(s([\mathbf{u}_1, \mathbf{c}])) + t([\mathbf{u}_1, \mathbf{c}]). \quad (6)$$

Since the subnetworks in a coupling block are never inverted themselves, we can append \mathbf{c} without losing the invertibility of the INN. Moreover, the operation has a lower-triangular jacobian whose determinant is the product of diagonal elements.

3.4.2 Dimensionality of latent. Central to an efficient style mapping framework is the dimensionality of the latent vector. Mapping the style of an image from the target domain to a low-dimensional style vector allows easy user-interactive image manipulation and editing. Our default architecture depicted in Fig. 2 encodes style in 3 latent dimensions, matching the dimensionality of an input pixel. However, some applications may require differently sized style vectors. As alternatives, we demonstrate the adaptations that enable the INN to operate with fewer or more latent dimensions.

To encode style in fewer dimensions, we split off some features after 4 invertible blocks. As depicted in Fig. 3 (left), these features are forced to follow the standard normal distribution using MLE. The remaining features continue through 4 more invertible blocks. The total number of invertible blocks is 8, to match our 3-dimensional INN from Fig. 2. This splitting of intermediate features is similar to the multiscale version of the normalizing flow described in [Dinh et al. 2016].

For a higher-dimensional latent style vector, we construct an invertible model on an augmented input space \mathbf{x}_{aug} . Similar to [Huang et al. 2020], each input vector \mathbf{x} is appended with samples from a standard normal distribution as shown in Fig. 3 (right). This allows us to improve the expressibility of the latent space at the cost of an increased dimension (from 3 to 4) for final image manipulation.

3.5 Optimization and inference

Similar to other normalizing flows, we train our INN by MLE where the likelihood is given by Eq. 5. The network thus learns to conditionally map a target pixel to a latent vector. However, when presented with an entire frame, there is no easy way to extract a single low-dimensional vector that captures style. For that, we augment MLE training with a reconstruction loss to force the per-pixel

representations of the same frame to lie closer in the latent space (see Fig. 1).

Likelihood loss: To restrict the magnitude of gradients for backpropagation, it is customary to minimize the negative log-likelihood (NLL) instead of likelihood since the logarithmic transform is monotonic. The NLL loss is:

$$\begin{aligned} \mathcal{L}_{\text{NLL}}(\mathbf{x}_p, \mathbf{y}_p) &= -\log p_Z(g_\theta^{-1}(\mathbf{x}_p; \mathbf{c}_p)) - \log |\nabla g_\theta^{-1}(\mathbf{x}_p; \mathbf{c}_p)| \\ &= \frac{\log 2\pi + (g_\theta^{-1}(\mathbf{x}_p; \mathbf{c}_p))^2}{2} - \log |\nabla g_\theta^{-1}(\mathbf{x}_p; \mathbf{c}_p)| \end{aligned} \quad (7)$$

This result follows since we choose a standard normal $\mathbf{z}_p \sim \mathcal{N}(0, 1)$ as the base distribution with the following log-likelihood:

$$\log p_Z(\mathbf{z}) = -\frac{1}{2} \log 2\pi - \frac{1}{2} \|\mathbf{z}\|_2^2. \quad (8)$$

The first term is a constant w.r.t. \mathbf{z} and can be dropped during training. Due to this loss, the INN learns a bijective mapping from the distribution of pixels to the chosen latent, conditioned on the encoded input pixel. Through MLE, we encourage per-pixel latents to follow a standard normal distribution as shown in Fig. 4 (right).

Reconstruction loss: Our secondary requirement is for pixels of the same frame to cluster together in the latent space Z . We achieve this by first passing all K pixels of frame i through the INN in reverse, and computing the centroid of their latent representations,

$$\mathbf{z}^i = \frac{1}{K} \sum_{p=1}^K g_\theta^{-1}(\mathbf{x}_p^i; \mathbf{C}(\mathbf{y}_p^i)), \quad (9)$$

where \mathbf{z}^i is the per-frame style vector. Then, we reconstruct the frame with K forward passes of the INN using the single extracted style latent \mathbf{z}^i but different conditional inputs. Finally, we compute the reconstruction loss \mathcal{L}_{rec} for each pixel of an image as follows:

$$\mathcal{L}_{\text{rec}}(\mathbf{x}_p^i, \mathbf{y}_p^i) = \|\mathbf{g}_\theta(\mathbf{z}^i; \mathbf{C}(\mathbf{y}_p^i)) - \mathbf{x}_p^i\|_2 \quad (10)$$

The reconstruction loss enforces a similar value of the style vector \mathbf{z}^i for the entire frame. We also observe that this simple constraint on the latent representation allows mapping of dissimilar frames in the input space to distant and distinct regions in the latent space (see Fig. 1).

Optimization of our INN is done bi-directionally with the total loss as the sum of the NLL loss \mathcal{L}_{NLL} and the reconstruction loss




Figure 3: Changes in INN architecture to decrease the latent style vector to 2 dimensions (left) or increase to 4 dimensions (right).

\mathcal{L}_{rec} . After successful training, the INN can be used to extract an overall per-frame *style* vector. This is achieved by running a reverse pass for each pixel of a given image and computing the centroid according to Eq. 9. In Fig. 4, we show the change in the latent space for the training samples due to the addition of \mathcal{L}_{rec} .

4 RESULTS

In this section, we evaluate the efficacy of our method in encoding the style of a target domain image into a low-dimensional latent space for the task of forward and inverse tone mapping. First, we compare our approach of conditioning color mapping on a style with the traditional approach of conditioning on the input image (Sec. 4.2, comparison with HDRNet). Then, we demonstrate that the existing combination of PCC with dimensionality reduction approaches (PCA and VAE) gives far inferior results as compared to our INN (Sec. 4.3). We report results in terms of PSNR (for RGB values) and FLIP [Andersson et al. 2020] here, and CIELab in the appendix. We choose these metrics because they are sensitive to color differences (unlike SSIM).

4.1 Datasets

We rely on two sources of SDR-HDR pairs. For images, we use Adobe-MIT 5K dataset [Bychkovsky et al. 2011]. Each RAW image in this dataset was tone-mapped (retouched) by 5 experts, who produced results in different styles. We report the results for expert C in the main paper, and for other expert in the supplementary materials. All images are rescaled to the height of 480 pixels for faster training and then split into a random 80/20% train/test sets. Due to the pixel-wise formulation of our INN, the same trained model can be employed on images of different resolutions at inference time.

Due to the lack of any publicly available manually color graded video datasets, we decoded 3 Blu-ray movies, namely “BBC Planet Earth II Episode 3 - Jungles”, “BBC Blue Planet II Episode 5 - Green Seas” and “The Lego Batman Movie”. 4K HDR content is often sold with two disks — one color graded for 4K HDR and another for 1080p SDR displays. We took advantage of that by extracting content from both disks. The frames from SDR and HDR streams were time-synchronized by finding the offset that maximized cross-correlation. Finally, the frames were manually inspected to ensure close correspondence. For good diversity, we construct a sequence from each video by collecting every 120th frame. For each video sequence, the first 80% of frames are used for training and the remaining 20% for testing. On average, we have 500-800 frames in the training set per movie. Each frame is rescaled to a resolution

of half HD (960×540). Both SDR and HDR RGB pixel values are display encoded (BT.2020 + PQ for HDR, BT.709 + sRGB for SDR).

Publicly available datasets proposed in methods like [Chen et al. 2021; Kim et al. 2019] do not include manual color grading, but instead rely on Youtube’s automatic HDR to SDR conversion process. The primary objective of our work is to model manual color grading, making such datasets unsuitable for the task of distilling color artist’s style.

4.2 Conditioning on style vs. image content

The central assumption of current deep learning tone mapping methods is that the right mapping can be found by analyzing global and local image features [Gharbi et al. 2017]. Based on that assumption, the existing methods employ large convolutional and/or fully connected networks operating on an entire image. In contrast, our INN operates at the pixel level to effectively distill the image-specific style vector without analyzing the image content. Here, we test which approach can better predict the results of manual tone-mapping and color grading. We train our INN separately on pairs of HDR-SDR frames taken from each movie from the movies dataset (see Sec. 4.1). Our pixel-wise training scheme allows us to operate on high-resolution images. Inference for each frame of half HD resolution (960×540) takes 0.025 seconds. Each model is trained for 80 epochs with an initial learning rate of $5e - 4$ with gradual learning rate scheduling. Additionally, for a more challenging setting, we train our INN on the Adobe-MIT 5K dataset [Bychkovsky et al. 2011].

As a representative example of existing learning-based tone mapping, we compare our results with HDRNet [Gharbi et al. 2017] (implemented in PyTorch [Ge 2021]), retrained on the same data as our method. The numerical results and quality metric distributions, shown in Fig. 6, demonstrate a dramatic improvement of 10-16 dB as compared to HDRNet. This shows that information about the style is necessary to faithfully reproduce manually color-graded or retouched images. Although HDRNet is a much larger network, consisting of 482K trainable parameters compared to 31K for our INN, it cannot infer the target image based on the source image alone. A few example images shown in Fig. 5 demonstrate that HDRNet fails to reproduce accurate color and tones of the target images. Similar to HDRNet, we see from Table 1 that other state-of-the-art image enhancement methods that do not model style struggle to reconstruct image-specific retouching with an accuracy higher than 25 dB. Please note that due to the novelty of our approach of encoding style as meta data, a direct comparison between




Figure 4: The figure shows the distribution of the latent color representation for different frames from the training set with (left) and without (right) the reconstruction loss \mathcal{L}_{rec} . We see that the latent vectors of the training set follow a normal distribution. Frames are taken from the BBC documentary Planet Earth “Jungles” and the INN trained to encode the style of frames to 2 dimensions is used for better visualization.




Figure 5: Qualitative comparisons with different methods on 3 datasets for the task of forward tone mapping. The target for MIT5k dataset is the expert retouched image. Additional results are provided in the appendix.

Table 1: Comparison of our INN with state-of-the-art methods that assume one-to-one mapping for Expert C from the MIT5k dataset [Bychkovsky et al. 2011]. The substantial improvement in performance clearly demonstrates how distilling style can lead to almost perfect reconstruction accuracy. The values for our and HDRNet methods were obtained using the same train/test splits, and the values from other methods are taken from the respective works.

	HDRNet [Gharbi et al. 2017]	UPE [Wang et al. 2019]	GleNet [Kim et al. 2020a]	3DLUT [Zeng et al. 2020]	StarEnh [Song et al. 2021]	Curl [Moran et al. 2021]	DPE [Chen et al. 2018]	CRSNet [He et al. 2020]	DLPF [Moran et al. 2020]	Ours
PSNR	22.49	23.24	25.88	24.92	25.46	24.04	23.76	24.23	23.93	39.22
Params.	482K	1M	-	<600K	14M	1.4M	3.34M	36K	1.8M	31.4K

our work and the different image enhancement methods is not fair. However, in Table 1, we show the results to delineate the need of our approach of style conditioning over conventional CNN based image-to-image mapping to achieve near perfect reconstruction.

Next, we train our INN-based mapping for the task of inverse tone mapping. While inverse tone-mapping often involves bit-depth

expansion and hallucination of over- and under-exposed pixels [Eilertsen et al. 2021], here we focus on the problem of learning global SDR→HDR color mapping. We use the same movie datasets as for the tone-mapping task but swap the source and target frames. The results shown in Fig. 7 demonstrate a substantial improvement of 10-12 dB over HDRNet.




Figure 6: Comparison of results on different datasets for the task of forward tone mapping. Note that our method achieves a substantial improvement in performance compared to other dimensionality reduction methods across datasets. The purple ‘+’ in the plots show the mean and black ‘x’ show the lowest 5th percentiles. Note that the y-axis for FLIP metric has been reversed. The results for Ours, PCA and VAE are reported for the model that encodes the style representation into 3 dimensional latent vector.




Figure 7: Comparison of results on different datasets for the task of inverse tone mapping on the movies dataset. To adapt PSNR to HDR images, we use perceptually uniform PU21 transform [Mantiuk and Azimi 2021]. The labels for the violin plots are consistent with Fig. 6. The results for Ours, PCA and VAE are reported for the model that encodes the style representation into 3 dimensional latent vector.

4.3 Other dimensionality reduction techniques

As explained in Sec. 3.1, the color mapping can be expressed as PCC (Eq. 2) and then the size of the style matrix can be reduced using standard dimensionality reduction methods, such as PCA or VAE. Here, we compare those standard approaches with our INN.

Principal component analysis: We ran PCA on training pairs to reduce the flattened *style* matrix M_{flat} into the required number of latent dimensions (2–4). During test time, \hat{M}_{flat} is reconstructed from the principal components and used to map the colors. We observe from Fig. 6 (forward tone mapping) and Fig. 7 (inverse tone mapping) that the strict linearity assumption of PCA results in poor performance. Qualitative result comparisons are provided in Fig. 5.

Variational autoencoder: Next, we replaced the linear projection with a deep auto-encoder architecture. Since we are interested in an interpretable latent space such as the one depicted in Fig. 8, we opted for a VAE where the latent follows a normal distribution [Kingma

and Welling 2013]. The input to the VAE is the original matrix of polynomial coefficients M . The goal is to train the VAE so that the matrix can be predicted from a low-dimensional latent vector. The training loss includes a reconstruction loss between predicted (after matrix multiplication using decoded matrix \hat{M}) and ground truth pixels and the evidence lower-bound. For a fair comparison with our method, we chose a fully connected network with approximately the same number of parameters. We empirically found the best results for a network with 6 hidden layers for the encoder and the decoder, with a total of 31K trainable parameters. The network was trained for 500 epochs with an initial learning rate of $5e-4$ with gradual learning rate scheduling. The weights given to the reconstruction and the evidence lower-bound were $\lambda = 1$, $\delta = 1e-3$, respectively. Although the VAE attains higher quality scores than PCA and its results are comparable to HDRNet (which does not use a style vector), VAE still performs much worse than our INN (see Fig. 6 and 7).




Figure 8: Example mapping obtained by manipulating a 2-dimensional *style* vector. Both dimensions control brightness and color temperature. Such style space enables assisted color grading, which mimics the range of styles found in the training image pairs.

5 APPLICATIONS

5.1 Assisted color grading

Color grading is a manual, labor-intensive process that requires a substantial set of skills. Our method can be used to partially automate this process. First, we ask the color artist to manually color grade N scenes, which we use to train our model. Then, we require the color artist to adjust only two or three parameters for the remaining scenes. Such adjustment is much easier than using color grading tools with dozens of different color adjustments. The added benefit of using our mapping is that the style is likely to be more consistent across the movie than if a manual color grading tool was used.

Since we do not have access to RAW video frames, typically used for color grading, we demonstrate this application using HDR frames from Blu-ray movies as input and the SDR frames as the color-graded target. Our results, from Sec. 4, have already demonstrated that our mapping can faithfully reproduce the SDR target frames. Figures 8 and 10 show example frames generated by adjusting a 2-dimensional style vector. Each frame comes from a different movie, for which a separate INN was trained. The style space allows for convenient exploration of tone and color adjustments that have been applied to previously color graded frames. The dimensions of the space are easy to interpret: they represent the change of color temperature, contrast and brightness. In a supplementary video, we demonstrate a mock interface of the real-time color grading tool.




Figure 9: Rate-distortion curves, comparing the learned INN to JPEG XT. HDR Image quality (y-axis) is measured by computing PSNR on PU21 encoded images [Mantiuk and Azimi 2021].

5.2 Transmission of SDR and HDR video content

The current practice is to encode and distribute SDR and HDR content separately (on Blu-ray or via streaming), which approximately doubles the required storage space. There exist methods for concurrent SDR+HDR image [Artusi et al. 2019] and video coding [Mantiuk et al. 2006], but they require transmitting a substantial amount of additional data. Here we show that our color mapping

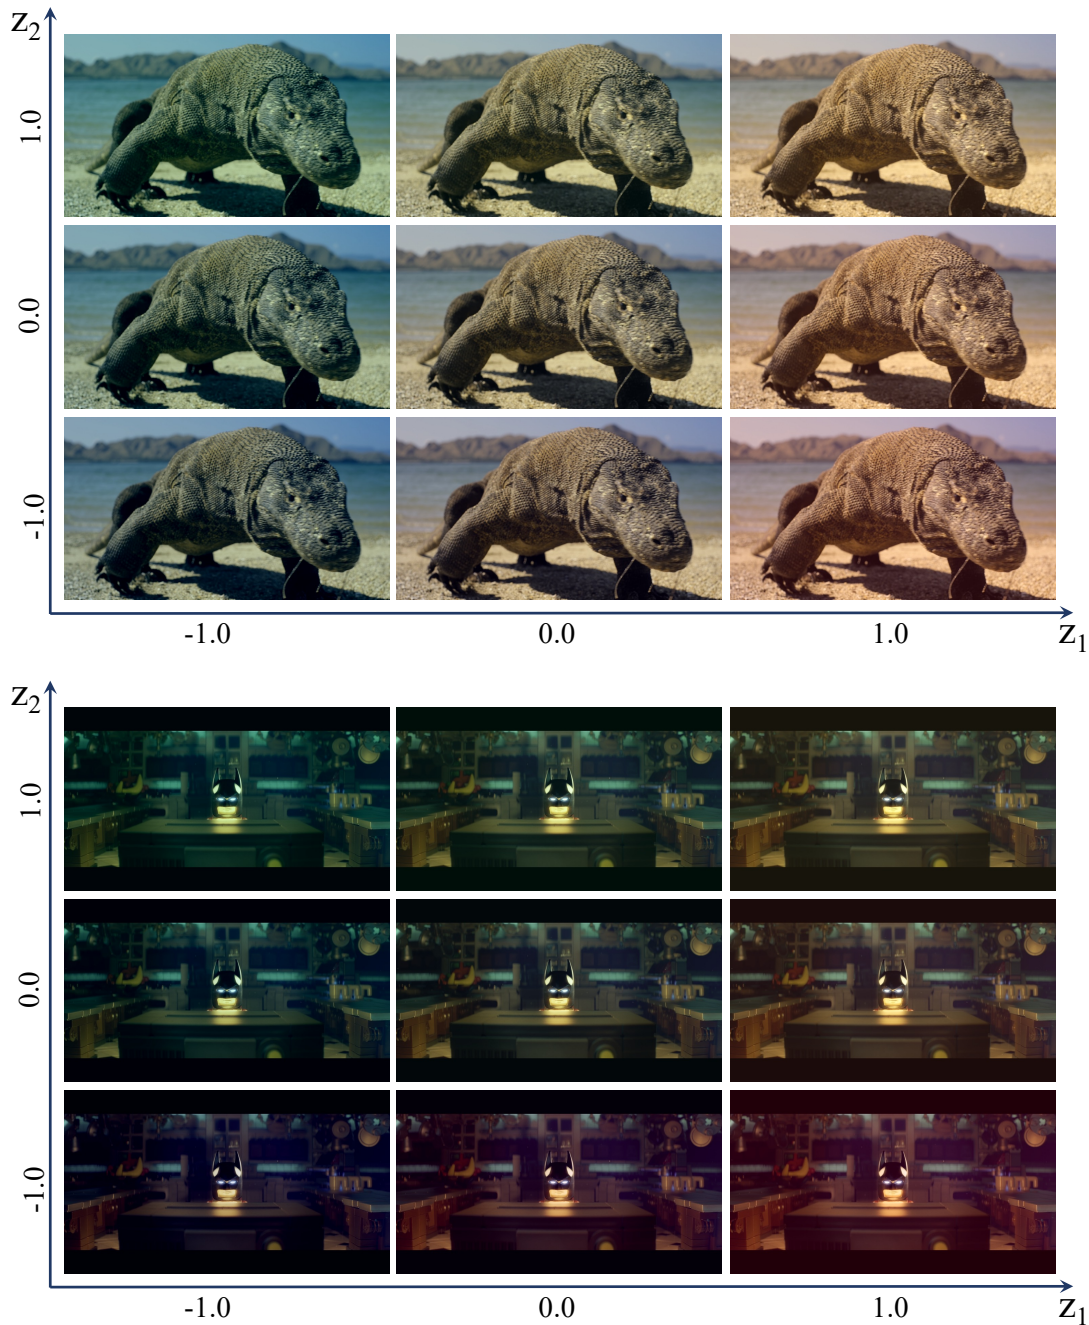


Figure 10: Additional results showing assisted color grading by manipulating a 2-dimensional *style* vector on two datasets - Planet Earth Episode 1 Islands (top) and Lego Batman movie (bottom). Both dimensions control brightness and color temperature.

can substantially reduce, or even eliminate, the need for auxiliary data.

We compare our inverse tone mapping INN from Sec. 4 with the coding used in JPEG XT (Profile A with open-loop encoding). The frames are encoded individually using either JPEG XT or a regular JPEG + our learned color mapping. JPEG XT encodes HDR images by storing a tone-mapped version of an HDR image (base layer), a

custom mapping function that predicts the HDR image from the tone-mapped image, and the difference between the predicted and the original HDR image (extension layer). Both base and extension layers are encoded using a standard JPEG codec. We replicate such encoding but replace the custom mapping used in JPEG XT with

	RGB	PCC-2	PCC-3	PCC-4	PCC-4 +VGG	PCC-4 +Hist	Dim-2	Dim-3	Dim-4
PSNR	28.98	31.52	33.02	41.68	30.130	35.69	17.67	18.24	18.35
len(c_p)	3	9	19	34	60	60	20.21	28.06	31.52
							37.85	41.68	41.99

Table 2: Ablation studies over the choice of conditioning vector (left) and over the dimensionality of the latent *style* vector for different methods (right). All results are reported for the “BBC Planet Earth Jungles” movie in terms of PSNR (dB).

our generative color mapping INN. Then, we measure the rate-distortion curves for a test set from the “Jungles” movie. The rate-distortion curves, shown in Fig. 9, depict a consistent improvement in performance when using our INN (employing both the base and extension layers). At an extremely low bit rate of 0.5, the INN without an extension layer produces the best quality images (PU-PSNR of over 30 dB).

5.3 Assisted dynamic range expansion for HDR displays

The vast majority of video content has been color graded for SDR displays and cannot take advantage of the higher luminance and contrast offered by HDR displays. Here, we show that it is possible to use our mapping function to expand SDR content for HDR displays. It should be noted, however, that our mapping will not be able to reconstruct details in the saturated parts of an image [Eilertsen et al. 2017a].

This process is identical to color grading, explained in the previous section, except we infer the HDR frames from their SDR counterpart. Similar to forward tone mapping, in this application we require only a small portion of the HDR frames to be manually color graded by color experts. Results of such assisted dynamic range expansion are included in the appendix.

6 ABLATION STUDY

To achieve the best performance, we conduct an ablation study over the choice of the conditioning vector c_p for training the INN. Furthermore, we provide a study over the dimensionality of the *style* vector.

6.1 Conditioning vector

First, we study the effect of using different degrees of polynomials in PCC as our per-pixel conditioning vector. We conducted an ablation study over the 1st (RGB), 2nd, 3rd and 4th degree polynomial expansion. Second, we study the effect of using image statistics as the conditioning vector in addition to the 4th degree polynomial. We train our INN architecture in such that each pixel-wise conditioning vector (PCC) is concatenated with high level features extracted from the entire image using an additional feed-forward network H . For this task, we use a pre-trained VGG network, similar to [Ardizzone et al. 2020; Denker et al. 2021] as the feed-forward network. The final conditioning vector for a given pixel is given as $c_p = [C(y_p), H(y)]$, where the weights of $H(\cdot)$ are simultaneously being updated alongside the weights of the INN. Finally, we train our INN using the 4th degree polynomial concatenated with the histogram of the luma channel as the conditioning vector. Table 2 (left) shows that the INN performs best when no additional image statistics are added to the 4th-degree PCC conditioning vector. Note

that different conditioning vectors have different lengths, as shown in the second row of Table 2 (left).

6.2 Dimensionality of the style vector

Next, we investigate the impact of changing the dimensionality of the latent *style* vector. Mapping the style of a target domain image to a low dimensional latent representation allows easy user-interactive image manipulation. To this end, we provide additional network architectures in Sec. 3.4.2 to train our INN for 2 and 4 latent style encodings. In Table 2 (right) we further provide a comparison of similar dimensions of the latent representations for PCA and VAE.

7 CONCLUSIONS

This work highlights the importance of modeling style when learning global image transforms, such as those between differently color-graded SDR-HDR images. We conclusively show that extracting a *style* vector from a target image considerably improves the reconstruction quality. This is due to the existence of an infinite number of equally plausible solutions, each representing a unique color artist’s choice. Our proposed conditional INN effectively models this one-to-many mapping by extracting and encoding this artistic choice from examples of image pairs into a low-dimensional *style* vector. We show that our method significantly outperforms state-of-the-art deep architectures that ignore style, as well as alternate dimensionality reduction methods that incorporate latent style but cannot encode it efficiently. Moreover, our invertible framework enables interactive style manipulation by adjusting the low-dimensional latent vector. The main focus of our work is color mapping for video content, which is global in nature (the same for all pixels in an image, $\mathcal{M} : \mathcal{R}^3 \rightarrow \mathcal{R}^3$). However, our method can be easily extended to a manual local mapping by splitting frames into a number of tiles, as done in [Eilertsen et al. 2015].

ACKNOWLEDGMENTS

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement N° 725253–EyeCode).

APPENDIX

In this appendix, we further report quantitative results for both forward and inverse tone-mapping on different datasets using additional metrics (Sec. 1). We also investigate whether a single INN can effectively capture all the expert styles from the MIT-Adobe FiveK dataset (Sec. 1.3). Finally in Sec. 2, we show improvement in performance due to the proposed bi-directional training with NLL \mathcal{L}_{NLL} and the reconstruction loss \mathcal{L}_{rec} .

1 ADDITIONAL RESULTS

1.1 Forward tone mapping

The violin plots in Fig. 11 compare our INN with HDRNet, PCA and VAE using the CIE DE 2000 [Sharma et al. 2005] metric. Similar to Fig. 6 in the main document, there is a substantial improvement in reconstruction quality due to correctly extracting and utilizing style.

For the MIT-Adobe FiveK dataset, we provide results for all the experts in Fig. 12. These are consistent with images retouched by expert C reported in the other figures. Here, a separate network is used for each expert to better compare with existing works.

1.2 Inverse tone mapping

For the task of inverse tone mapping, Fig. 13 shows similar violin plots for 2 more metrics: FLIP and CIE DE 2000. Before running the SDR metrics, we encode the reconstructed and ground truth HDR images with PU21 encoding.

1.3 Image content vs. style conditioning

Additionally, we trained a single INN for all 5 experts of MIT-Adobe FiveK, something that can not be done with HDRNet because of the lack of conditioning on style. Since we learn a one-to-many mapping, the same network produces outputs in different styles by utilising different latent vectors. Fig. 14 shows that our single INN successfully captures the style of all the experts. We further report the quantitative comparison of our single trained INN on the individual test set for each expert in Table 4. Refer to Fig. 15 and Fig. 16 for qualitative comparisons on selected scenes. The deep-learning methods, like HDRNet, that model expert retouching with one-to-one mappings are unsuitable for this task since a single network cannot learn different styles corresponding to the experts. We see that our INN performs much better than PCA and VAE and produces artifact-free images that better match the required style.

2 BI-DIRECTIONAL TRAINING

When trained with the NLL loss \mathcal{L}_{NLL} , our INN learns to conditionally map a target pixel to a latent vector. However, when presented with an entire frame, there is no easy way to extract a single low-dimensional vector, that captures the style of the mapped frames. Recall that we augment MLE training by forcing the per-pixel representations of the same frame to lie closer in the latent space.

In Table 3, we show the effect of our proposed bi-directional training, by addition of the reconstruction loss \mathcal{L}_{rec} alongside \mathcal{L}_{NLL} . The NLL loss makes the style vectors resemble a predetermined latent distribution (the standard normal in our experiments), while

the reconstruction loss ensures that pixels from the same frame have similar style vectors.

Table 3: Ablation study on the effect of bi-directional training for the task of forward tone-mapping for the “BBC Planet Earth II Episode 3 - Jungles” dataset.

	\mathcal{L}_{NLL}	$\mathcal{L}_{\text{NLL}} + \mathcal{L}_{\text{rec}}$
PSNR \uparrow	36.80	41.68
FLIP \downarrow	0.108	0.070

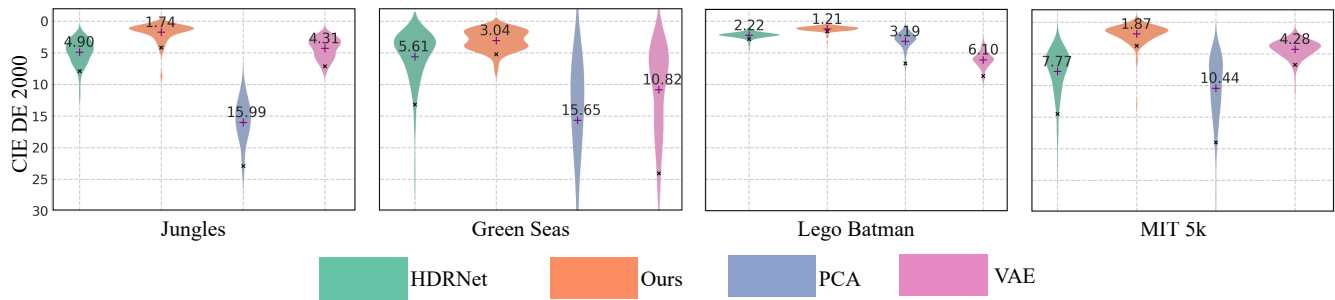


Figure 11: Comparison of results on different datasets for the task of forward tone mapping for CIE DE 2000 metric. Our method achieves a substantial improvement in performance compared to other dimensionality reduction methods across datasets. The purple '+' in the plots show the mean and black 'x' show the lowest 5th percentiles. Note that the y-axis for CIE DE 2000 metric has been reversed.

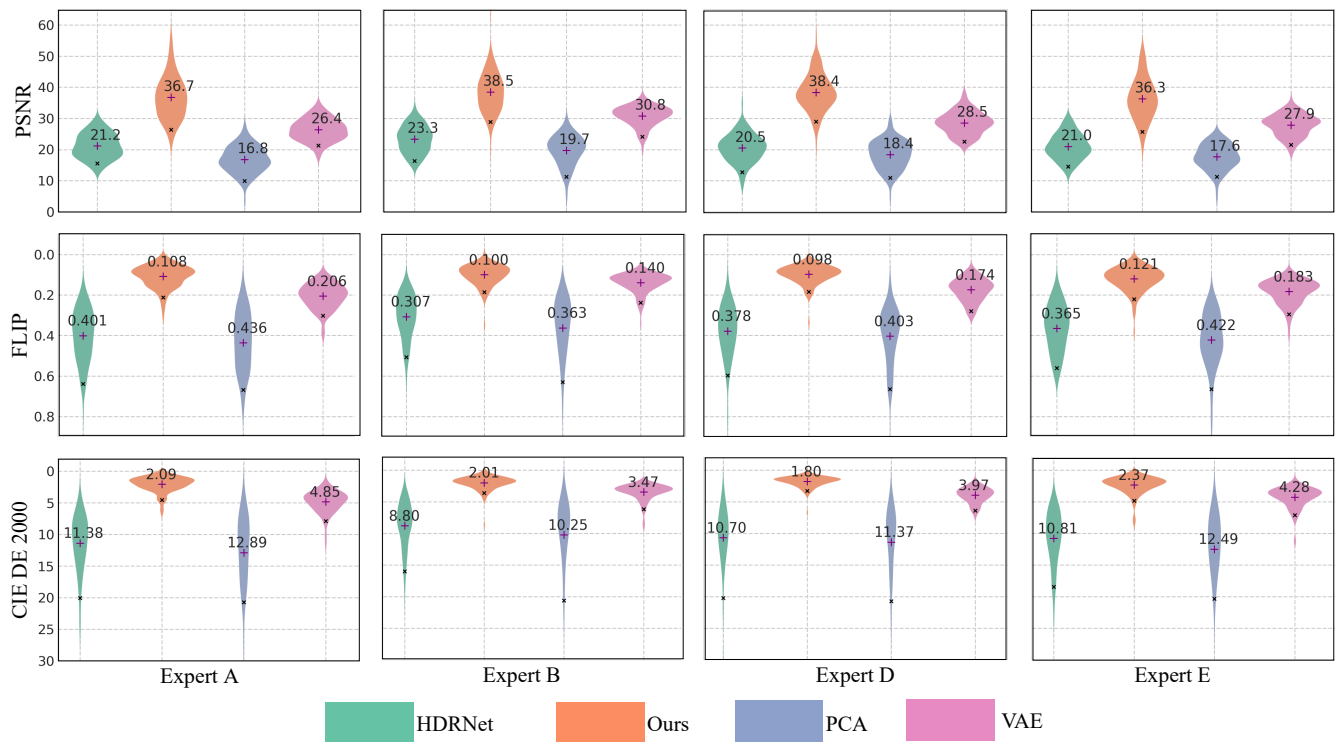


Figure 12: Comparison of our results with other methods on the remaining 4 experts from the MIT-Adobe FiveK dataset. Note that our method achieves a substantial improvement in performance compared to other dimensionality reduction methods across experts. The purple '+' in the plots show the mean and black 'x' show the lowest 5th percentiles. Note that the y-axis for FLIP and CIE DE 2000 metric have been reversed.

Table 4: Additional quantitative results for each expert from the MIT-Adobe FiveK dataset for a single model trained on all experts. Results are reported in terms of PSNR, FLIP and CIE DE 2000 metric. Note that HDRNet is unsuitable for this task since a single network cannot learn different styles corresponding to the experts.

	Expert A	Expert B	Expert C	Expert D	Expert E
PSNR \uparrow	36.4	38.1	38.9	38.2	36.0
FLIP \downarrow	0.116	0.104	0.099	0.100	0.124
CIE DE \downarrow	2.11	2.02	1.91	1.93	2.38




Figure 13: Here we show the results for the task of inverse tone mapping on two other metrics: FLIP and CIE DE 2000. To adapt the metrics to HDR images, we use perceptually uniform PU21 transform [Mantiuk and Azimi 2021]. The purple ‘+’ in the plots show the mean and black ‘x’ show the lowest 5th percentiles. Note that the y-axis for FLIP and CIE DE 2000 metric have been reversed.




Figure 14: Here we show a comparison of our method with PCA and VAE when a single model is trained for all experts from the MIT-Adobe FiveK dataset. We report the results in terms of PSNR, FLIP and CIE DE 2000 metric. Note that the y-axis for FLIP and CIE DE 2000 metric have been reversed. The test set in this case contains 5000 images, 1000 for each expert. Additional inference results for each expert individually are reported in Table 4.




Figure 15: Qualitative comparison of methods that simultaneously capture the different styles of all experts of the MIT-Adobe FiveK dataset. The bottom row shows the reference images. Quantitative results over all test images are plotted in Fig. 14. Note that HDRNet is unsuitable for this task since a single network cannot learn different styles corresponding to the experts.

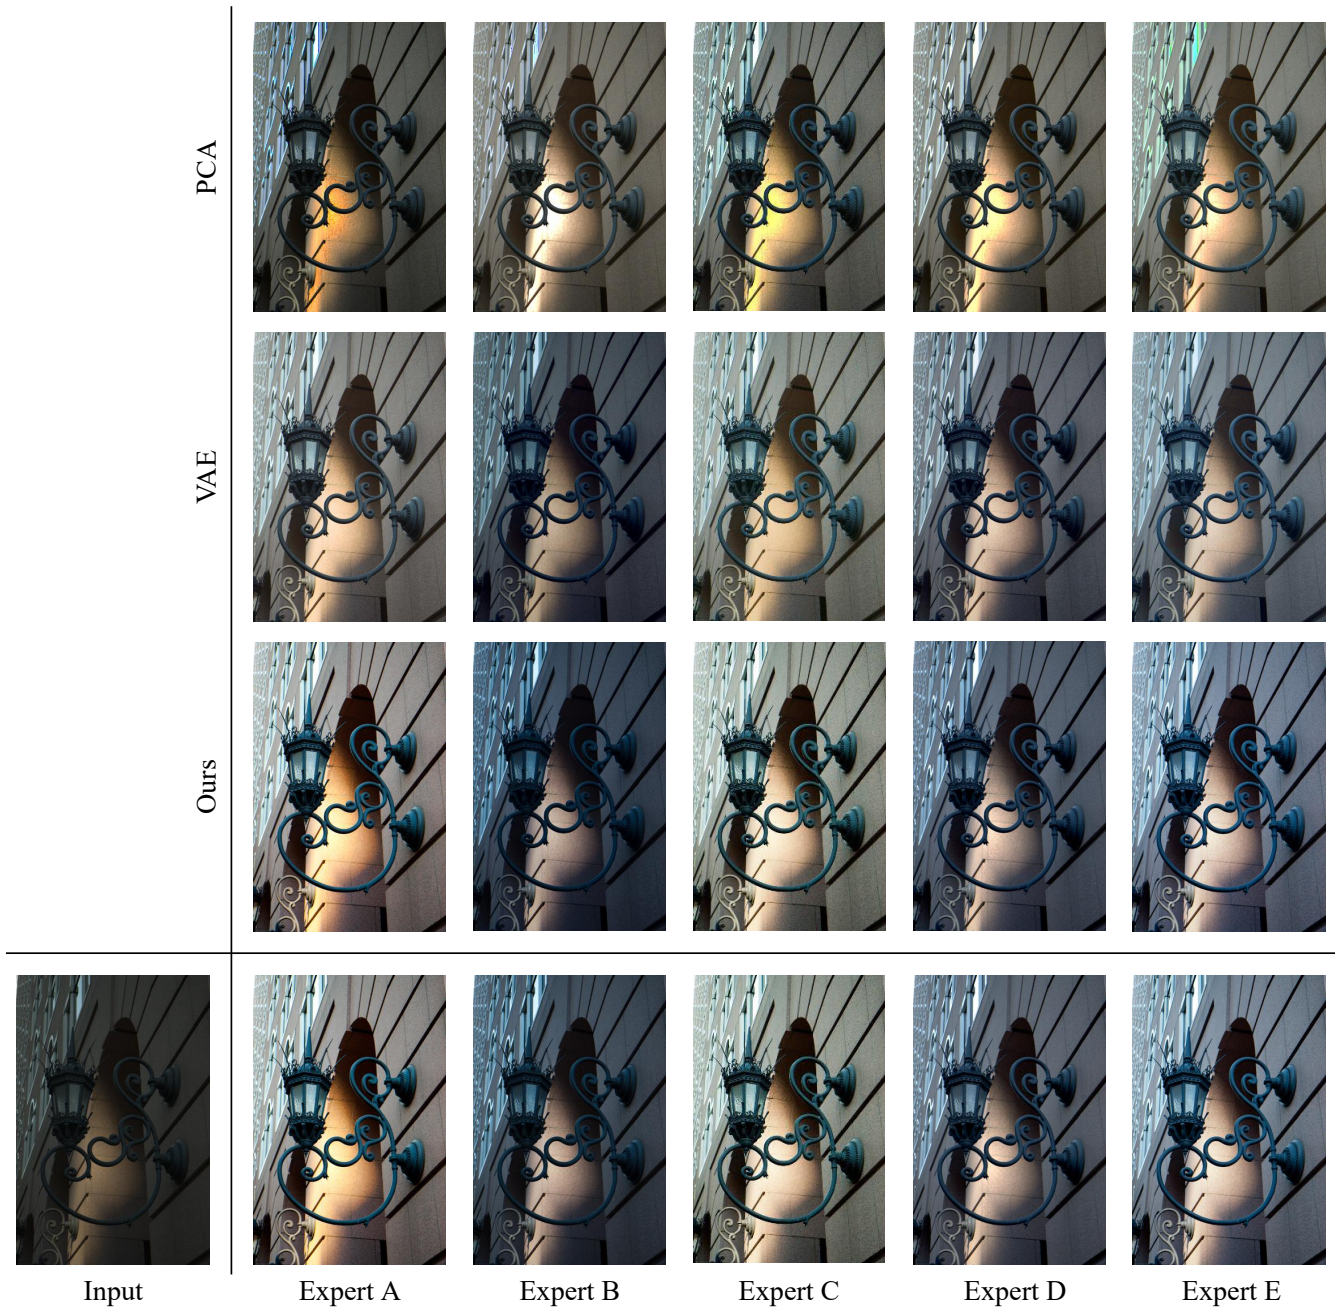


Figure 16: Here we show qualitative comparison of methods that simultaneously capture the different styles of all experts of the MIT-Adobe FiveK dataset for a different image. The bottom row shows the reference images.

REFERENCES

- Pontus Andersson, Jim Nilsson, Tomas Akenine-Möller, Magnus Oskarsson, Kalle Åström, and Mark D. Fairchild. 2020. FLIP: A Difference Evaluator for Alternating Images. *Proc. ACM Comput. Graph. Interact. Tech.* 3, 2, Article 15 (aug 2020), 23 pages. <https://doi.org/10.1145/3406183>
- Lynton Ardiszone, Jakob Kruse, Carsten Lüth, Niels Bracher, Carsten Rother, and Ullrich Köthe. 2020. Conditional invertible neural networks for diverse image-to-image translation. In *DAGM German Conference on Pattern Recognition*. Springer, 373–387.
- Alessandro Artusi, Rafal K. Mantiuk, Thomas Richter, Philippe Hanhart, Pavel Korshunov, Massimiliano Agostinelli, Arkady Ten, and Touradj Ebrahimi. 2019. Overview and evaluation of the JPEG XT HDR image compression standard. *Journal of Real-Time Image Processing* 16, 2 (Apr 2019), 413–428. <https://doi.org/10.1007/s11554-015-0547-x>
- Vladimir Bychkovsky, Sylvain Paris, Eric Chan, and Fredo Durand. 2011. Learning photographic global tonal adjustment with a database of input/output image pairs. In *CVPR 2011*. IEEE, 97–104. <https://doi.org/10.1109/CVPR.2011.5995413>
- Xiangyu Chen, Zhengwen Zhang, Jimmy S. Ren, Lynhoo Tian, Yu Qiao, and Chao Dong. 2021. A New Journey from SDRTV to HDRTV. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Yu-Sheng Chen, Yu-Ching Wang, Man-Hsin Kao, and Yung-Yu Chuang. 2018. Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6306–6314.
- Alexander Denker, Maximilian Schmidt, Johannes Leuschner, and Peter Maass. 2021. Conditional Invertible Neural Networks for Medical Imaging. *Journal of Imaging* 7, 11 (2021), 243.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. 2016. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803* (2016).
- Gabriel Eilertsen, Saghi Hajisharif, Param Hanji, Apostolia Tsirikoglou, Rafal K Mantiuk, and Jonas Unger. 2021. How to cheat with metrics in single-image HDR reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3998–4007.
- Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafal K. Mantiuk, and Jonas Unger. 2017a. HDR image reconstruction from a single exposure using deep CNNs. *ACM Transaction on Graphics* 36, 6 (2017), Article 178.
- Gabriel Eilertsen, Rafal K. Mantiuk, and Jonas Unger. 2015. Real-time noise-aware tone mapping. *ACM Transactions on Graphics* 34, 6 (oct 2015), 1–15. <https://doi.org/10.1145/2816795.2818092>
- G. Eilertsen, R. K. Mantiuk, and J. Unger. 2017b. A comparative review of tone-mapping algorithms for high dynamic range video. *Computer Graphics Forum* 36, 2 (May 2017), 565–592. <https://doi.org/10.1111/cgff.13148>
- Graham D Finlayson and Mark S Drew. 1997. Constrained least-squares regression in color spaces. *Journal of electronic imaging* 6, 4 (1997), 484–493.
- Graham D Finlayson, Michal Mackiewicz, and Anya Hurlbert. 2015. Color correction using root-polynomial regression. *IEEE Transactions on Image Processing* 24, 5 (2015), 1460–1470.
- Jinchen Ge. 2021. HDRnet-PyTorch. <https://github.com/gejinchen/HDRnet-PyTorch>.
- Michaël Gharbi, Jiawen Chen, Jonathan T Barron, Samuel W Hasinoff, and Frédo Durand. 2017. Deep bilateral learning for real-time image enhancement. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 118.
- Jingwen He, Yihao Liu, Yu Qiao, and Chao Dong. 2020. Conditional sequential modulation for efficient global image retouching. In *European Conference on Computer Vision*. Springer, 679–695.
- Chin-Wei Huang, Laurent Dinh, and Aaron Courville. 2020. Augmented normalizing flows: Bridging the gap between generative flows and latent variable models. *arXiv preprint arXiv:2002.07101* (2020).
- Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. 2017. Dslr-quality photos on mobile devices with deep convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*. 3277–3285.
- Hanul Kim, Su-Min Choi, Chang-Su Kim, and Yeong Jun Koh. 2021. Representative Color Transform for Image Enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4459–4468.
- Han-Ul Kim, Young Jun Koh, and Chang-Su Kim. 2020a. Global and local enhancement networks for paired and unpaired image enhancement. In *European Conference on Computer Vision*. Springer, 339–354.
- Han-Ul Kim, Young Jun Koh, and Chang-Su Kim. 2020b. PieNet: Personalized Image Enhancement. In *European Conference on Computer Vision*.
- Soo Ye Kim, Jihyong Oh, and Munchul Kim. 2019. Deep SR-ITM: Joint Learning of Super-Resolution and Inverse Tone-Mapping for 4K UHD HDR Applications. In *Proceedings of the IEEE International Conference on Computer Vision*.
- Durk P Kingma and Prafulla Dhariwal. 2018. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems* 31 (2018).
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- Yu-Lun Liu, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. 2020. Single-image HDR reconstruction by learning to reverse the camera pipeline. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'20)*. 1651–1660.
- Andreas Lugmayr, Martin Danelljan, Luc Van Gool, and Radu Timofte. 2020. Srflow: Learning the super-resolution space with normalizing flow. In *European conference on computer vision*. Springer, 715–732.
- Rafal Mantiuk, Scott Daly, and Louis Kerofsky. 2008. Display Adaptive Tone Mapping. *ACM Trans. Graph.* 27, 3 (aug 2008), 1–10. <https://doi.org/10.1145/1360612.1360667>
- R Mantiuk, A Efreinov, K Myszkowski, and H.-P. Seidel. 2006. Backward compatible high dynamic range MPEG video compression. In *ACM Transactions on Graphics*. <https://doi.org/10.1145/1141911.1141946>
- Rafal K. Mantiuk and Maryam Azimi. 2021. PU21: A novel perceptually uniform encoding for adapting existing quality metrics for HDR. *Picture Coding Symposium (PCS)* (2021), 1–5. <https://doi.org/10.1109/PCS50896.2021.9477471>
- Demetris Marnerides, Thomas Bashford-Rogers, Jonathan Hatchett, and Kurt Debatista. 2018. Expandnet: A deep convolutional neural network for high dynamic range expansion from low dynamic range content. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 37–49.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *European conference on computer vision*. Springer, 405–421. https://doi.org/10.1007/978-3-030-58452-8_24
- Sean Moran, Pierre Marza, Steven McDonagh, Sarah Parisot, and Gregory Slabaugh. 2020. DeepLFP: Deep Local Parametric Filters for Image Enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sean Moran, Steven McDonagh, and Gregory Slabaugh. 2021. CURL: Neural Curve Layers for Global Image Enhancement. In *2020 25th International Conference on Pattern Recognition (ICPR)*. 9796–9803. <https://doi.org/10.1109/ICPR48806.2021.9412677>
- Jongchan Park, Joon-Young Lee, Donggeun Yoo, and In So Kweon. 2018. Distort-and-recover: Color enhancement using deep reinforcement learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5928–5936.
- Aakanksha Rana, Praveer Singh, Giuseppe Valenzise, Frederic Dufaux, Nikos Komodakis, and Aljosa Smolic. 2020. Deep Tone Mapping Operator for High Dynamic Range Images. *IEEE Transactions on Image Processing* 29 (2020), 1285–1298. <https://doi.org/10.1109/TIP.2019.2936649>
- Erik Reinhard, Michael Stark, Peter Shirley, and James Ferwerda. 2002. Photographic tone reproduction for digital images. *ACM Transactions on Graphics* 21, 3 (jul 2002), 267. <https://doi.org/10.1145/566654.566575>
- Danilo Rezende and Shakir Mohamed. 2015. Variational inference with normalizing flows. In *International conference on machine learning*. PMLR, 1530–1538.
- Marcel Santana Santos, Tsang Ing Ren, and Nima Khademi Kalantari. 2020. Single image HDR reconstruction using a CNN with masked features and perceptual loss. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 80–1. <https://doi.org/10.1145/3386569.3392403>
- Gaurav Sharma, Wencheng Wu, and Edul N Dalal. 2005. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application: Endorsed by Inter-Society Color Council, The Colour Group (Great Britain), Canadian Society for Color, Color Science Association of Japan, Dutch Society for the Study of Color, The Swedish Colour Centre Foundation, Colour Society of Australia, Centre Français de la Couleur* 30, 1 (2005), 21–30.
- Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- Yuda Song, Hui Qian, and Xin Du. 2021. StarEnhancer: Learning Real-Time and Style-Aware Image Enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4126–4135.
- TG Stockham Jr. 1972. Image processing in the context of a visual model. *Proc. IEEE* 60, 7 (1972), 828–842. <https://doi.org/10.1109/PROC.1972.8782>
- Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. 2020. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems* 33 (2020), 7537–7547.
- J. Tumblin and H. Rushmeier. 1993. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications* 13, 6 (1993), 42–48. <https://doi.org/10.1109/38.252554>
- Ruixing Wang, Qing Zhang, Chi-Wing Fu, Xiaoyong Shen, Wei-Shi Zheng, and Jiaya Jia. 2019. Underexposed photo enhancement using deep illumination estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6849–6857.
- Zhicheng Yan, Hao Zhang, Baoyuan Wang, Sylvain Paris, and Yizhou Yu. 2016. Automatic Photo Adjustment Using Deep Neural Networks. *ACM Transactions on Graphics* 35, 2 (May 2016), 1–15. <https://doi.org/10.1145/2790296>
- Akiko Yoshida, Rafal Mantiuk, Karol Myszkowski, and Hans-Peter Seidel. 2006. Analysis of Reproducing Real-World Appearance on Displays of Varying Dynamic Range. *Computer Graphics Forum (Proc. of Eurographics)* 25, 3 (Sep 2006), 415–426. <https://doi.org/10.1111/j.1467-8659.2006.00961.x>

Hui Zeng, Jianrui Cai, Lida Li, Zisheng Cao, and Lei Zhang. 2020. Learning image-adaptive 3d lookup tables for high performance photo enhancement in real-time.

IEEE Transactions on Pattern Analysis and Machine Intelligence (2020).