

Flexible Example-based Image Enhancement with Task Adaptive Global Feature Self-Guided Network

Dario Kneubuehler, Shuhang Gu, Luc Van Gool, and Radu Timofte

Computer Vision Lab, ETH Zurich, Switzerland

Abstract. We propose the first practical multitask image enhancement network, that is able to learn one-to-many and many-to-one image mappings. We show that our model outperforms the current state of the art in learning a single enhancement mapping, while having significantly fewer parameters than its competitors. Furthermore, the model achieves even higher performance on learning multiple mappings simultaneously, by taking advantage of shared representations. Our network is based on the recently proposed SGN architecture, with modifications targeted at incorporating global features and style adaption. Finally, we present an unpaired learning method for multitask image enhancement, that is based on generative adversarial networks (GANs).

1 Introduction

Digital images are omnipresent in today’s society, with a wide scope of applications ranging from posting snapshots taken with smartphones on social media, to high profile fashion shoots and photojournalism. Current image enhancement software provides tools to locally and globally adjust images to one’s liking, yet the use of such tools requires a considerable amount of time and the results highly depend on the user’s skills. Automating this task through the use of algorithms remains an important and challenging task to this day. In this paper, we study the problem of automatic image enhancement through deep neural networks (DNN). Image enhancement comprises a wide set of different image processing tasks ranging from image denoising, super resolution to illumination adjustment. We focus on the task of example-based image enhancement, which aims to enhance a group of low-quality images to the quality of another group of high quality images. Early image enhancement works [2,5,17] mainly focused on contrast enhancement and illumination estimation using algorithmic models and traditional machine learning methods. Recently, the success of deep neural networks (DNN) on other computer vision tasks triggered the study of DNN-based image enhancement approaches. Compared with traditional heuristic or prior based approaches, DNN methods that directly learn a mapping function from the low quality image to high quality image, have achieved promising enhancement performance. Despite their impressive performance, current DNN-based methods are limited by the following three issues.

Network architecture. Compared with classical image restoration tasks such as image denoising and super-resolution, the image illumination enhancement tasks requires global adjustment of the input and thus a very large receptive field is inevitable. Although several network architectures have been proposed for the image illumination enhancement task, better architectures are necessary to achieve a better trade-off between performance and efficiency.

Training data. In addition to enhancing the image texture locally, enhancement algorithms also need to adjust the global image illumination, therefore, a large number of training samples is required to provide image-level supervision. However, the acquisition of paired data for image enhancement is laborious and costly. In [3], Bychkovsky *et al.* provide a dataset with 5000 raw input images and corresponding professionally retouched versions, that is well suited for use with supervised machine learning methods. However, 5000 pairs of images are far from enough to train a good enhancement network. In [12], Ignatov *et al.* collected weakly paired data from different devices, *e.g.*, cell phones and DSLR camera. However, the authors only provide roughly aligned patches (100×100 px) which significantly limits the valid receptive field of an enhancement network in the training phase.

Limited flexibility. The working scenarios of enhancement algorithms are very complex, which requires the enhancement network to be highly flexible. On the one hand, low quality images produced by different devices, or by the same device in different environmental conditions have distinct characteristics. The poor generalization of existing DNN-based algorithms requires us to train different networks for different types of degradation, which is highly impractical. On the other hand, image enhancement is a highly subjective task, for the same input image, different people might favor different enhancement results. Adapting to the preference of different customers is very important.

In this paper, we propose an Global feature Self-Guided Network (GSGN) to address all the above issues. To efficiently exploit large scale contextual information, we modify the recently proposed self-guided network [8] by adding multiple modification targeted at incorporating global features. Specifically, we follow the top-down strategy of SGN and incorporate the large scale contextual information at an early stage to guide the following processing steps, by adding a global feature branch on top of it. Besides incorporating global information at an early stage via the global feature branch, we constantly incorporate image scale information using instance normalization (IN) layers on all levels of the network. Concretely, for handling different enhancement tasks with a compact model, we switch the instance normalization layers with adaptive instance normalization Layers (AdaIN) and add a small auxiliary network that transforms the latent input. Using this scheme we train the network for different tasks and only allow the AdaIN layers to be task-specific. Another advantage of such a multi-task learning strategy lies in a data augmentation perspective. As the network is able to take advantage of shared feature representations and leverage training data from multiple tasks, our method achieves better results on both the many-to-one

mapping and one-to-many mapping cases. In summary, our main contributions are:

- We propose a novel SGN [8] based deep neural network architecture that outperforms its competitors by a large margin on the task of supervised image to image mapping, while requiring considerably fewer parameters.
- We show that by using multitask learning to learn multiple mappings simultaneously, we can take advantage of shared representations which yields even higher performance compared to learning one-to-one mappings separately.
- For the supervised and unsupervised settings, we conduct experiments for both, one-to-many and many-to-one mappings. The experimental results validate the effectiveness of multi-task learning for image enhancement.

2 Related Work

Image enhancement is a classical yet hot topic in the area of computer vision. While initial research in the field was based on algorithms that rely on heuristic rules, such as histogram equalization and retinex-based methods, recent research has shifted to learning based methods. Specially, as a powerful tool for image to image mapping, deep neural networks (DNNs) have achieved great success in image enhancement. Previous works have investigated the DNN-based enhancement methods from different aspects.

One category of studies aims to investigate better network for capture the mapping function between input and target image. Yan *et al.* [18] trained a multilayer feed forward neural network to capture the mapping function between two groups of images. Lore *et al.* [14] used an autoencoder based network to tackle low-light image enhancement. Gharbi *et al.* [6] presented a bilateral learning approach that is optimized for real time performance on smartphones. Chen *et al.* [19] augmented the U-Net approach [15] with a global branch to better capture the global information. Wang *et al.* [16] presented a supervised learning method for improving underexposed photos that relies on an intermediate illumination mapping.

Another category of research attempts to push image enhancement toward real application scenarios. To improve the quality of cell phone images, Ignatov *et al.* [12] collected images with different devices and roughly aligned image patches for training enhancement network. However, as accurate image registration is a challenging problem, Ignatov *et al.* [12] only provide image patches for training, which greatly limited the capacity of network to leverage global information. To enhance images for real applications, unsupervised learning approaches have been proposed. Deng *et al.* [4] used adversarial learning for aesthetic-driven image enhancement with weak supervision. Chen *et al.* [19] proposed a deep photo enhancer which relies on a two-way generative adversarial network (GAN) architecture.

While some of the recently proposed models show impressive results for supervised and weakly supervised image enhancement, none of them is able to

learn more than a single deterministic mapping function. However, users of enhancement software are known for having different aesthetic preferences, that cannot be satisfied by providing them with a single mapping function. This limitation is also non-ideal from the perspective of neural network training, since a lot more data would be available if the models were able to learn multiple tasks simultaneously. The next logical step is to introduce multi-task learning into image enhancement and provide the user with high level controls that can satisfy multiple stylistic preferences.

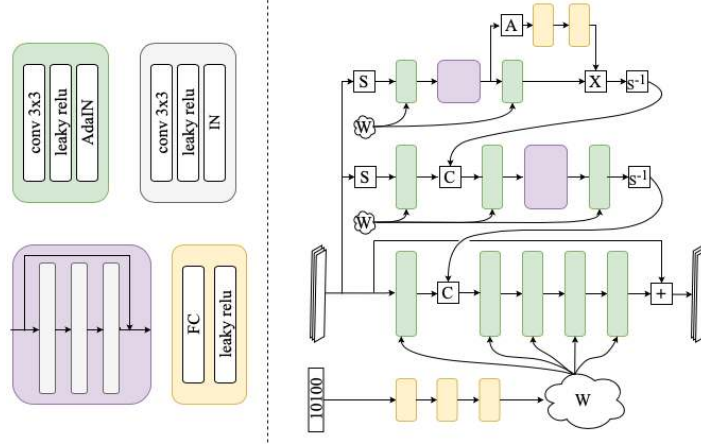


Fig. 1. Architecture of MT-GSGN network. The latent vector z is mapped to a intermediate latent space w , which is used as input to the adaptive instance normalization layers (AdaIn). For GSGN no mapping network is used and AdaIn layers are replaced with instance normalization (IN) layers. S stands for the shuffling operation, S^{-1} for the inverse shuffling operation, and C for concatenation. A denotes global average pooling. Residual blocks, convolutional blocks and fully connected blocks, have purple, green and orange color coding respectively.

3 Proposed Methods

In this section, we introduce our Global feature Self-Guided Network (GSGN) model and how it can be extend to the multi-tasks version, *i.e.* MT-GSGN. First, we briefly review the network architecture of the original SGN model. Then, we introduce modifications targeted at solving the illumination enhancement problem. Lastly, we present how the task-specific global feature branch can be utilized, and propose the flexible MT-GSGN method.

3.1 Brief introduction to SGN [8]

Gu *et al.* [8] proposed the Self-Guided Network (SGN) to more efficiently incorporate large scale contextual information for image denoising. In order to have an

overview of the image content from large receptive field, the SGN method adopts a top-down guidance strategy. Specifically, shuffling operations are adopted to generate multi-resolution inputs, and SGN firstly processes the top-branch and gradually propagates the features extracted at low spatial resolution to higher resolution sub-networks. With the effective self-guided mechanism, the SGN [8] has achieved state-of-the-art denoising performance.

3.2 Global Feature Self-Guided Network (GSGN)

As SGN has been shown to be an effective network architecture for incorporating large scale contextual information, we adopt it to solve the image illumination enhancement task, for which a large receptive field is even more essential. Although the receptive field of SGN is much larger than the other denoising networks, it is still insufficient for the illumination enhancement task. Furthermore, SGN has a considerable number of parameters, training it with the limited number of samples in the enhancement datasets might lead to over-fitting during training or unused networks capacity. In order to adapt the network to the enhancement task, we introduce the Global feature Self-Guided Network (GSGN). The architecture of the proposed GSGN is shown in Fig. 1. GSGN differs from SGN in the following aspects.

Global Feature Incorporation. We incorporate global features by performing global average pooling of the feature maps in the top most branch of GSGN and feeding it through two fully connected layers. These global features are then multiplied with the output feature maps on the same branch. In contrast to the global feature scheme used in [19], our approach works for arbitrary size input images.

Less Parameters. In order to reduce the number of parameters, we reduce the number of levels from 3 to 2 and reduce the number of channels in the higher level sub-networks by a factor of two. At the same time we increase the number of feature maps in the base level sub-network by a factor of two.

Constantly Incorporating Global Information. In spite of incorporating the global information in an early stage of the network, we employ instance normalization (IN) layers after the activation functions of convolutional layers. While instance normalization is mostly used as a means for stabilizing training, we use it to learn global features in all levels of the network. Each instance normalization layer has two parameters for each channel. These parameters are learned dependent on the feature maps of a single image and therefore act as global features. As we have deployed complex global transforms in an earlier stage, such simple IN layer helps us to constantly incorporate global information.

With the above modifications we obtain a relatively lightweight GSGN model with only 339k parameters. In section 5.1, we provide an ablation study to validate our design choices.

3.3 Flexible Image Enhancement with Task Adaptive GSGN

To enable task adaptive learning, GSGN is augmented with an additional mapping network that consists of multiple fully connected layers with leaky relu activation functions, similar to the technique used in [13]. The number of fully connected layers was empirically set to 3. The mapping network takes a latent vector z as input, which encodes the desired style to be learned, and maps it to an intermediate latent space w . This transformed latent vector is then used as the input to adaptive instance normalization [11] layers that are inserted after the convolution layers. Furthermore, an additional fully connected layer is used before each adaptive instance normalization layer to match the dimension of w to the dimension of the relevant feature map.

3.4 Unsupervised Flexible Enhancement with Task Adaptive GSGN

Two way GAN architectures have proven to be a suitable solution for unsupervised image to image mapping problems. To show the suitability of our Task Adaptive GSGN for unsupervised learning, we use the popular CycleGAN architecture [21], and add an additional network that provides a conditional loss based on the task. An illustration of our unpaired training setup is shown in figure 2. In general, the setup can be described as a two player game where player one, called the generator G tries to conditionally produce fake samples X'_t given a prior sample X_s . Player two, called the critic D tries to evaluate how close the fake samples X'_t are to the real samples X_t . In our case X_s are image samples from the source distribution, while X_t are high quality samples from the target distribution and X'_t is the estimated enhanced version of the input image. In addition to the generator $G_{s \rightarrow t}$ that maps samples from the source domain to the target domain, a CycleGAN architecture uses an additional generator $G_{t \rightarrow s}$ that maps samples from the target domain back to the source domain. Using these mappings, we define two cyclic mappings between the domains: 1) $X_s \rightarrow X'_t \rightarrow X''_s$ and 2) $X_t \rightarrow X'_s \rightarrow X''_t$. Since both $G_{s \rightarrow t}$ and $G_{t \rightarrow s}$ are optimized using the same loss function, we only use the term G in reference for both networks.

Adversarial Loss We use a version of the WGAN-GP loss [9], where a critic is used instead of a discriminator. In contrast to loss functions based on the original GAN formulation [7], the output of the network is not feed through a sigmoid function but used directly to approximate the wasserstein distance between two probability distributions. In order for this to work the network needs to be constrained to a 1-Lipschitz function. This is achieved by putting a penalty on the gradients of the critic during training. The gradient penalty λ is computed by evaluating the gradients of linearly interpolated samples \hat{y} between X_t and X'_t :

$$\lambda = \max(0, \|\nabla \mathbb{D}(\hat{y})\|_2 - 1). \quad (1)$$

While, the loss functions for the critic D amounts to:

$$L_D = (D(X_t) - D(X'_t))\lambda w, \quad (2)$$

where w is a hyper-parameter used to control the amplitude of the gradient penalty λ . The adversarial loss for the generator G can be computed as:

$$L_A = D(X'_t). \quad (3)$$

Cycle Loss To ensure cycle consistency of the mappings between distributions, we impose a cycle loss between X_s and X_s'' as well as between X_t and X_t'' . Defined as:

$$L_{CY} = MSE(X_s, X_s'') + MSE(X_t, X_t'') \quad (4)$$

Illumination invariant identity loss The use of identity losses between the source and target images is common in CycleGAN setups, to ensure that the content of the processed image is similar to the input image. However, since the change in illumination between input and output is large for the image enhancement task, traditional MSE and MAE pixel based losses result in high identity losses for target mapping. Because illumination enhancement results in a shift of the mean pixel values, a simple measure to mitigate this problem is to subtract the mean over the pixel values of each image before computing the identity loss. Using the this definition, the identity loss can be stated as follows:

$$L_I = MSE(X_s - \mu(X_s), X'_t - \mu(X_t)) + MSE(X_t - \mu(X_t), X'_s - \mu(X'_s)), \quad (5)$$

where μ is the arithmetic mean function, applied over all the pixel values of an image.

Conditional Loss For the one-to-many case, we use an additional network C that acts as a classifier on the images produced and provides a loss based on reconstructing the latent vector z from the generated image. Therefore we define the conditional loss:

$$L_{CO} = -(z \log(C(X'_t)) + (1 - z) \log(1 - C(X'_t))). \quad (6)$$

The network C is optimized using the same loss function but computed over X_t instead of X'_t .

Total Loss The total loss for the Generator network G amounts to:

$$L_G = L_{CY}w_{CY} + L_Iw_I + L_Aw_A + L_{CO}w_{CO} \quad (7)$$

Where w_{CY} , w_I , w_A , and w_{CO} are weights to balance the contribution of the different losses to the total generator loss.

4 Experimental Setup

4.1 Datasets

MIT5K [3]. The dataset is composed of 5,000 high resolution images that are retouched by five experts performing global and local adjustments. We follow the experimental setting of [19] and use 2,250 images and their retouched adaptations for training the supervised models. The test set contains 500 images.

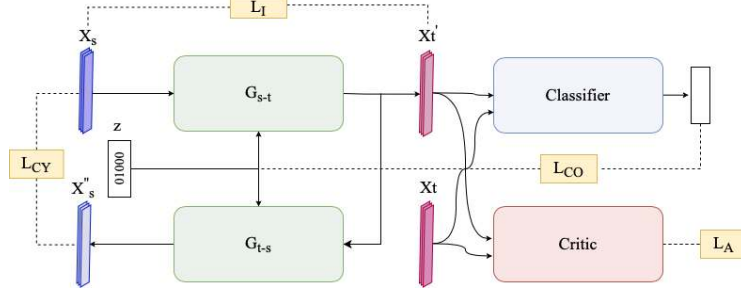


Fig. 2. Two way GAN training setup used for learning one-to-many mappings with unpaired data. Note that for simplicity only the cycle loss from X_s to X_s'' is displayed, while the second cycle from X_t to X_t'' is omitted

The remaining 2,250 images are reserved for the target domain in the unsupervised setting. For the multitask experiments data of all five experts is used, thus increasing the amount of train and test images to 11,250 and 2,500, respectively.

DPED [12]. DPED contains paired data of scenes captures with three smartphone cameras, *i.e.*, iPhone 3GS, BlackBerry Passport and Sony Xperia Z, and a professional quality DSLR, Canon 70D. In total the dataset contains 22K full resolution photos. To get paired samples suited for supervised training, the authors generated aligned patches of 100×100 pixels. There are 139K, 160K and 162K training patches for BlackBerry, iPhone and Sony smartphones respectively. Correspondingly, there are 2.4K, 4.4K and 2.5K test patches.

Flickr Multi-style. To test our approach on a dataset in the wild, we created a small multi class dataset. There are three style classes. 1) *Normal*, which consists of a subset of 600 images from the MIT5K dataset. 2) *Sunset*, which consists of 600 images with the tag sunset that are collected from Flickr and selected according to Flickr’s interestness score. 3) *HDR*, which is a subset of 594 of the HDR-Flickr dataset from [19] that are also collected using the Flickr API. Some visual examples of the constructed dataset can be found in Fig. 3.



Fig. 3. Visual examples of the images in different subset of the constructed Flickr dataset.

4.2 Performance measures

In our experiments we use the standard Peak Signal to Noise Ratio (PSNR) and Structural Similarity Index (SSIM) to measure the fidelity of the enhanced images towards the ground truth / target reference images. Complementary we use LPIPS [20], a learned perceptual metric meant to approximate the human perceptual similarity judgements.

4.3 Implementation Details

We use the TensorFlow framework for all experiments. For the supervised experiments on the MIT5K dataset [3], we follow the setup used in [19]. This includes using the exact train and test split, in order to get comparable results. The size of the input layer is fixed at 512×512 pixels during training. Images are resized for the longer edge to equal 512 pixels prior to training. To match the input of the network the images are zero padded. In all supervised experiments we use the following cost function, which maximizes the PSNR: $L_G = -\log_{10}(MSE(I_s, I_t))$. The models are trained for 100k iterations with batch size 4 and learning rate $1e-4$ on MIT5K dataset, while on DPED dataset the models are trained for 150k iterations with batch size 50 and learning rate $1e-4$.

5 Results on Supervised Image Enhancement

In this section, we evaluate our method for supervised image enhancement. For more visual results we refer to the supplementary material.

5.1 Single task enhancement

To justify our design choices as well as to compare with other state-of-the-art approaches, we firstly conduct experiments on the supervised single task enhancement setting. We follow the experimental setting of [19] and train different models to approximate the retouched results by expert C in the MIT5K dataset [3].

Concretely, we compare the proposed GSGN network with the original 2 levels SGN [8], denoted by SGN2; and the SGN2 with reduced number filter channel numbers, denoted by GSGN w/o global feat. & instance normalization; SGN2 with reduced number filter channel numbers and global feature branch, denoted by GSGN w/o instance normalization; and our final GSGN model. In addition to our ablation study, we also compare GSGN with three recently proposed architectures that were designed for this task. The comparison methods include DPED [12], DPE [19] and UPEDIE [16]. DPED [12] uses a fully convolutional

| Method | PSNR | SSIM | Para. |
|----------------------------|--------------|--------------|-------|
| DPED [12] | 21.76 | 0.871 | 401k |
| DPE [19] | 23.80 | 0.900 | 5019k |
| UPEDIE [16] | 23.04 | 0.893 | - |
| SGN2 | 20.98 | 0.863 | 769k |
| GSGN w/o global feat. & IN | 21.19 | 0.865 | 325k |
| GSGN w/o IN | 23.74 | 0.900 | 338k |
| GSGN (Ours) | 24.16 | 0.905 | 339k |

Table 1. Comparison between state-of-the-art and our GSGN with different configurations on MIT5K dataset (Expert C). We refer the reader to the latest **corrected** version of the UPEDIE paper [16].



Fig. 4. Visual comparison of learned supervised mappings of our GSGN network against the DPE [19] method on the test set of the MIT5K dataset.

ResNet for image enhancement. DPE [19] used U-Net [15] as a backbone and augmented it with global features. UPEDIE [16] adopts an alternative approach by learning an illumination map from an encoder network. Table 1 shows the average PSNR and SSIM of the different approaches, evaluated on the 500 test images of the MIT5K dataset. Figure 4 shows visual results of the proposed GSGN network. For the results reported in [19] we were not able to reproduce their high PSNR and SSIM, to be fair we still use the original numbers they claim in the paper. The results show, that our network not only outperforms its competitors by at least 0.36 dB, but also uses considerably less parameters to achieve these results. The results in table 1 validate the effectiveness of the proposed GSGN architecture. GSGN achieved higher PSNR and SSIM while using far fewer parameters than current state-of-the-art methods.

5.2 Multi-task Enhancement: One-to-Many

The fact that the Adobe MIT-5K dataset contains five retouched versions of each input image, makes it ideally suited to be used as a base for multitask learning in a supervised one-to-many setting. In order to demonstrate the feasibility of our task adaptive GSGN model, we train models to capture all five input to output mappings of the MIT5K dataset simultaneously. We consider the following three experimental settings. Results are found in table 2.

Single Task Models. We train our GSGN model on the task of learning all input to output mappings separately for all five experts A to E in a traditional one-to-one setting. The averaged PSNR over all five experiments is 23.97 dB. It is worth noting that there is quite a large difference of up to 3.82 dB PSNR between the different experts. This gives interesting insights into the dataset. On one side

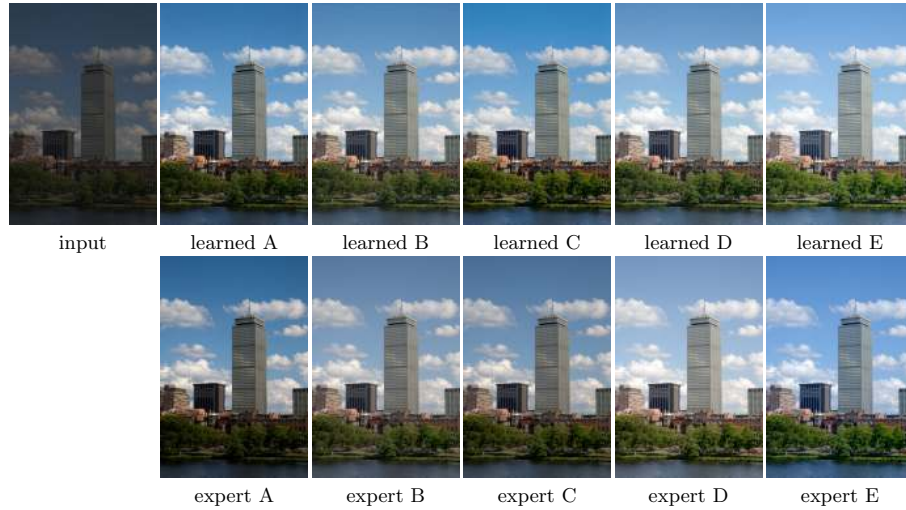


Fig. 5. Comparison of learned supervised mappings by our MT-GSGN model and corresponding expert labels.

this could mean that certain experts were more consistent when applying the mapping. On the other hand, this could also indicate that some experts used more complex adjustments than others, which made it harder for the network to learn these mappings.

Single Model for All Tasks. The GSGN model is trained on data of all five task simultaneously in a traditional one-to-many setting but without a constant latent embedding so it can not explicitly discriminate between the different tasks. This effectively increases the amount of training data by a factor of five. However, the network is only able to learn an average mapping and cannot take into account the specific style of each expert. While performance gains come from being trained on larger amounts of data, the lack of being able to differentiate between the different experts causes the performance to decrease, this results in a lower average PSNR of 23.30 dB.

Multi-tasking with a Task Adaptive Model. We use our task adaptive GSGN model to learn all five tasks simultaneously, in a one-to-many setting. This combines the advantages of the previous two settings. The resulting model reaches the highest average PSNR of 24.32 dB and also outperforms the other settings in terms of SSIM, MSSIM and LPIPS [20]. This not only shows that the model is able to efficiently use the labels to learn multiple distributions simultaneously, but also that by taking advantage of the larger amounts of data available in a multitask setting, the model is able to outperform models that were just trained for a single task. Figure 5 shows visual results of the MT-GSGN model and the corresponding expert labels.

| Tasks \ Methods | Task-Specific Models | | | Single Model for All Tasks | | | Multi-tasking (MT-GSGN) | | |
|--------------------------|----------------------|-------|--------|----------------------------|-------|-------|-------------------------|--------------|--------------|
| | PSNR | SSIM | LPIPS | PSNR | SSIM | LPIPS | PSNR | SSIM | LPIPS |
| Supervised enhancement | | | | | | | | | |
| Raw → A | 22.30 | 0.876 | 0.0743 | 21.28 | 0.862 | 0.080 | 22.51 | 0.879 | 0.072 |
| Raw → B | 26.12 | 0.949 | 0.0443 | 25.48 | 0.939 | 0.049 | 26.53 | 0.952 | 0.043 |
| Raw → C | 24.16 | 0.905 | 0.0610 | 23.71 | 0.896 | 0.069 | 24.44 | 0.904 | 0.063 |
| Raw → D | 23.02 | 0.903 | 0.0621 | 23.14 | 0.905 | 0.063 | 23.27 | 0.906 | 0.060 |
| Raw → E | 24.23 | 0.924 | 0.0571 | 22.92 | 0.911 | 0.067 | 24.86 | 0.928 | 0.055 |
| Avg. | 23.97 | 0.911 | 0.0598 | 23.30 | 0.903 | 0.066 | 24.32 | 0.914 | 0.059 |
| Unsupervised enhancement | | | | | | | | | |
| Raw → A | 18.95 | 0.771 | 0.212 | 16.63 | 0.705 | 0.289 | 19.99 | 0.824 | 0.123 |
| Raw → B | 20.49 | 0.840 | 0.204 | 18.23 | 0.809 | 0.255 | 22.30 | 0.908 | 0.097 |
| Raw → C | 19.16 | 0.779 | 0.228 | 16.76 | 0.707 | 0.304 | 20.81 | 0.839 | 0.120 |
| Raw → D | 17.92 | 0.767 | 0.246 | 16.46 | 0.743 | 0.287 | 20.36 | 0.866 | 0.111 |
| Raw → E | 17.75 | 0.758 | 0.271 | 16.03 | 0.730 | 0.311 | 19.89 | 0.859 | 0.124 |
| Avg. | 18.85 | 0.783 | 0.232 | 16.82 | 0.739 | 0.289 | 20.67 | 0.859 | 0.115 |

Table 2. Enhancement one-to-one and **one-to-many** results on MIT5K.

5.3 Multi-task Enhancement: Many-to-One

We further validate our model on a many-to-one multi-task mapping on DPED [12] dataset and train models to enhance the images from the iPhone, Sony and Blackberry cameras to the DSLR camera domain. We follow the same three step approach described in the previous section, and compare the average performance of the single task model against the multi-task model. It is worth noting that in this case a weighted average is needed because the number of images in the training sets for each phone is different.

| Method | Supervised | | | Unsupervised | | |
|--------------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | PSNR | SSIM | LPIPS | PSNR | SSIM | LPIPS |
| iPhone | 22.94 | 0.819 | 0.142 | 19.47 | 0.648 | 0.264 |
| Sony | 24.46 | 0.877 | 0.103 | 23.21 | 0.821 | 0.142 |
| Blackberry | 23.18 | 0.842 | 0.108 | 19.44 | 0.640 | 0.225 |
| Weighted avg. | 23.41 | 0.841 | 0.122 | 20.48 | 0.640 | 0.225 |
| Single Model (all) | 23.35 | 0.836 | 0.129 | 20.22 | 0.685 | 0.238 |
| Multi-tasking | 23.69 | 0.839 | 0.128 | 20.69 | 0.749 | 0.209 |

Table 3. Enhancement one-to-one and **many-to-one** results on DPED.

Table 3 shows the evaluation results for this experiment. In this case the multi-task model outperforms its single task counterparts, in terms of PSNR, but not in terms of SSIM and LPIPS. This is likely due to the fact that the model is trained to optimize PSNR but training patches coming from different sources can never be perfectly aligned due to nonlinear distortions coming from different lenses and sensors, as well as perspective distortions. The authors of [12] note that there could be shift of up to 5 pixels between source and target images. This makes the problem formulation ill posed for supervised learning

and suggests that this dataset is better suited for weakly supervised learning methods.

6 Results on Unsupervised Image Enhancement

Implementation Details As described in section 3, we follow a two cycle GAN approach inspired by [1], where an additional conditional loss is used to provide task specific gradients during training. A discriminator architecture similar to the one described in [19] is used. We train the critic 30 times more than the generator for the MIT5K and Flickr Multi-style datasets and 40 times more often for the DPED dataset.

Multi-task Enhancement: One-to-Many We evaluate our model on MIT5K dataset, using our weakly supervised GAN architecture. Analog to the supervised experiments, we perform the same three experimental settings. Table 2 shows the validation results. Our multi-task learning approach with task adaptive GSGN again outperforms, our GSGN network by taking advantage of shared feature representations.

Multi-task Enhancement: Many-to-One We validate the proposed algorithm on the many-to-one task on the DPED dataset. In contrast to the one-to-many setting, this setting does not require a conditional loss on the generated images, which simplifies the training. Table 3 shows the evaluation results.

Flickr Multi-style Dataset Finally, we evaluate our MT-GSGN model on a dataset on the wild. For this purpose we train it on our Flickr dataset on the task of learning a one-to-many mapping with three styles to learn, normal, sunset and hdr. Figure 6 shows visual examples of the learned mappings. Note that the model is also able to interpolate between styles, even though it was not explicitly trained for this task. Figure 7 shows a visual comparison of the learned mappings against the state of the art. For more visual results we refer to the supplementary material.

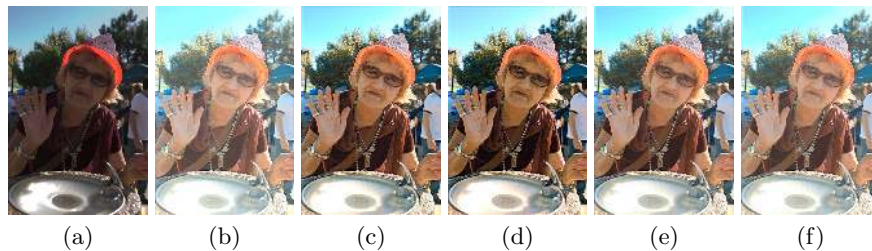


Fig. 6. Visual examples of enhanced images from Task adaptive GSGN trained on the Flickr dataset. (a) Input, (b) Normal, (c) Interpolation between HDR and Sunset, (d) HDR, (e) Sunset, (f) Interpolation between Normal and HDR. We refer to sec. 6 for details.

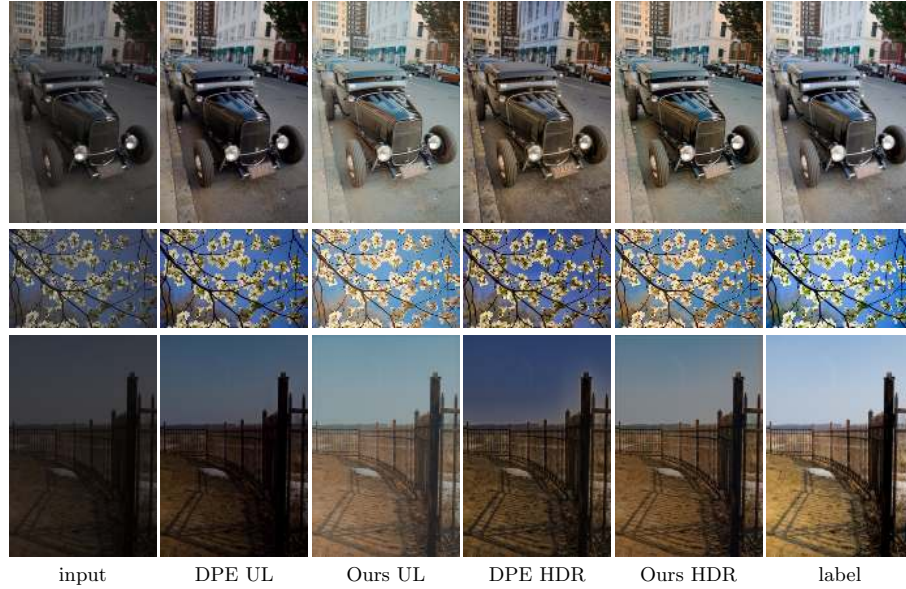


Fig. 7. Visual comparison of learned unsupervised mappings by our Task Adaptive GSGN model trained on the small Flickr Multi-style dataset, against the DPE [19] method. Note that our model learns both mappings in a single training while DPE requires separate networks and multiple trainings.

7 Conclusion

In this paper, we have investigated and proposed a flexible example-based image enhancement network, that uses a task adaptive global feature self-guided network. First, we proposed a novel network architecture capable to outperform existing methods on the task of supervised image to image mapping while requiring much fewer parameters. Second, we demonstrated that by using multi-task learning we benefit from shared representation and achieve higher performance compared to learning one to one mappings separately. Third, for both supervised and unsupervised settings our experimental results validate the effectiveness of our multi-task learning for image enhancement in one-to-many and many-to-one settings. To the best of our knowledge this is the first successful work in applying multi task learning to the challenging image enhancement problem.

Acknowledgements

This work was partly supported by ETH Zurich General Fund (OK), by a Huawei project and by Amazon AWS and Nvidia grants.

References

1. Almahairi, A., Rajeswar, S., Sordoni, A., Bachman, P., Courville, A.C.: Augmented cyclegan: Learning many-to-many mappings from unpaired data. CoRR **abs/1802.10151** (2018), <http://arxiv.org/abs/1802.10151>
2. Aubry, M., Paris, S., Hasinoff, S.W., Kautz, J., Durand, F.: Fast local laplacian filters: Theory and applications. ACM Transactions on Graphics (TOG) **33**(5), 167 (2014)
3. Bychkovsky, V., Paris, S., Chan, E., Durand, F.: Learning photographic global tonal adjustment with a database of input / output image pairs. In: The Twenty-Fourth IEEE Conference on Computer Vision and Pattern Recognition. Colorado Springs, CO (June 2011)
4. Deng, Y., Loy, C.C., Tang, X.: Aesthetic-driven image enhancement by adversarial learning. CoRR **abs/1707.05251** (2017), <http://arxiv.org/abs/1707.05251>
5. Farbman, Z., Fattal, R., Lischinski, D., Szeliski, R.: Edge-preserving decompositions for multi-scale tone and detail manipulation. In: ACM Transactions on Graphics (TOG). vol. 27, p. 67. ACM (2008)
6. Gharbi, M., Chen, J., Barron, J.T., Hasinoff, S.W., Durand, F.: Deep bilateral learning for real-time image enhancement. ACM Transactions on Graphics (TOG) **36**(4), 118 (2017)
7. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks (2014)
8. Gu, S., Li, Y., Gool, L.V., Timofte, R.: Self-guided network for fast image denoising. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
9. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of wasserstein gans. CoRR **abs/1704.00028** (2017), <http://arxiv.org/abs/1704.00028>
10. Hu, Y., He, H., Xu, C., Wang, B., Lin, S.: Exposure: A white-box photo post-processing framework. CoRR **abs/1709.09602** (2017), <http://arxiv.org/abs/1709.09602>
11. Huang, X., Belongie, S.J.: Arbitrary style transfer in real-time with adaptive instance normalization. CoRR **abs/1703.06868** (2017), <http://arxiv.org/abs/1703.06868>
12. Ignatov, A., Kobyshev, N., Timofte, R., Vanhoey, K., Van Gool, L.: Dslr-quality photos on mobile devices with deep convolutional networks. In: the IEEE Int. Conf. on Computer Vision (ICCV) (2017)
13. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. CoRR **abs/1812.04948** (2018), <http://arxiv.org/abs/1812.04948>
14. Lore, K.G., Akintayo, A., Sarkar, S.: Llnet: A deep autoencoder approach to natural low-light image enhancement. CoRR **abs/1511.03995** (2015), <http://arxiv.org/abs/1511.03995>
15. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
16. Wang, R., Zhang, Q., Fu, C.W., Shen, X., Zheng, W.S., Jia, J.: Underexposed photo enhancement using deep illumination estimation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)

17. Wang, S., Zheng, J., Hu, H.M., Li, B.: Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE Transactions on Image Processing* **22**(9), 3538–3548 (2013)
18. Yan, Z., Zhang, H., Wang, B., Paris, S., Yu, Y.: Automatic photo adjustment using deep learning. *CoRR* **abs/1412.7725** (2014), <http://arxiv.org/abs/1412.7725>
19. Yu-Sheng Chen, Yu-Ching Wang, M.H.K., Chuang, Y.Y.: Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In: *CVPR* (2018)
20. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. *CoRR* **abs/1801.03924** (2018), <http://arxiv.org/abs/1801.03924>
21. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks (2017)

A Flexible Example-based Image Enhancement with Task Adaptive Global Feature Self-Guided Network – Supplementary

A.1 Supervised learning - Visual comparison against the State-of-the-art

Figures A1 and A2 show a visual comparison of learned supervised mappings by our GSGN model against the state-of-the-art method DPE [19] on additional images from the test split of the MIT5K [3] dataset. While the DPE model is able to approximate the target in many cases, our model yields more consistent results and is able to handle a larger variety of input images.

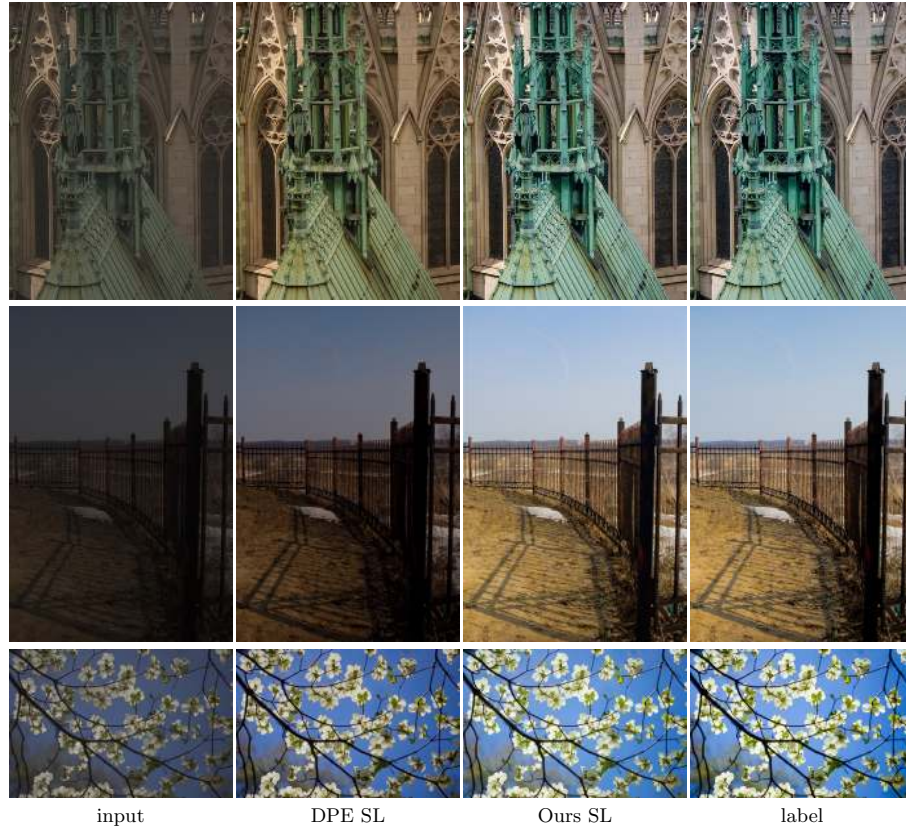


Fig. A1. Visual comparison of learned supervised mappings of our GSGN network against the DPE [19] method on images of the test set of the MIT5K [3] dataset (expert C).

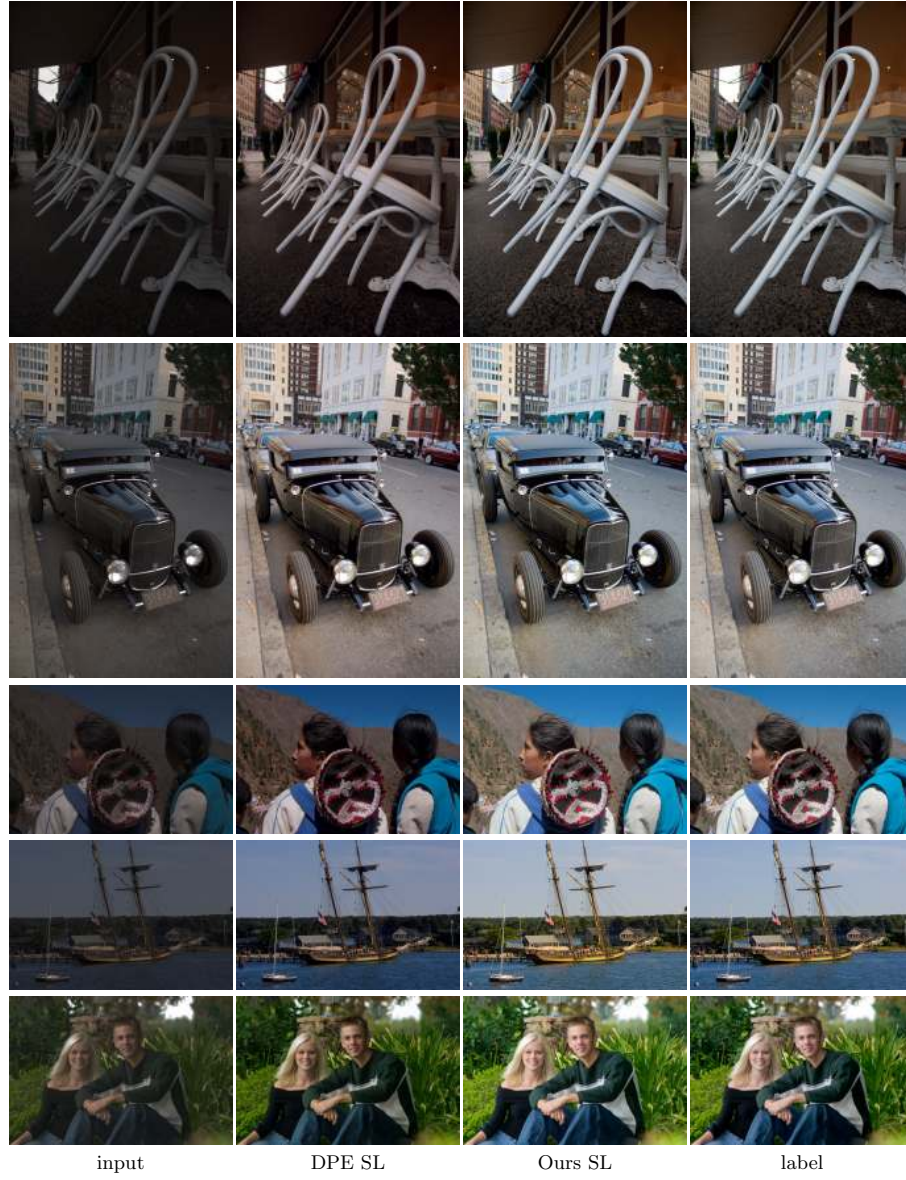


Fig. A2. Visual comparison of learned supervised mappings of our GSGN network against the DPE method on images of the test set of the MIT5K dataset (expert C).

A.2 Supervised learning - Visual results of multitask enhancement experiment

Figure A3 and figure A4 show results of our MT-GSGN model on the task of supervised multitask learning on the five experts of the MIT5K [3] dataset. The metrics in table 2 in the main paper show impressive results of the models ability to learn the distribution of the five experts. However, looking at single samples shows little variation between experts in both training data and predictions made by the model. Considering that the human retouches of the different experts have limited consistency, this might also mean that a significant portion of an experts style is random. However, our result still has relevance. Even though the variance between experts on single images is limited, computing the average distance over the whole test set (500 images per expert) in terms of PSNR, SSIM and LPIPS [20] (table 2 in the paper) shows consistent results and proves that non random aspects of the different expert distributions can be learned by our model.

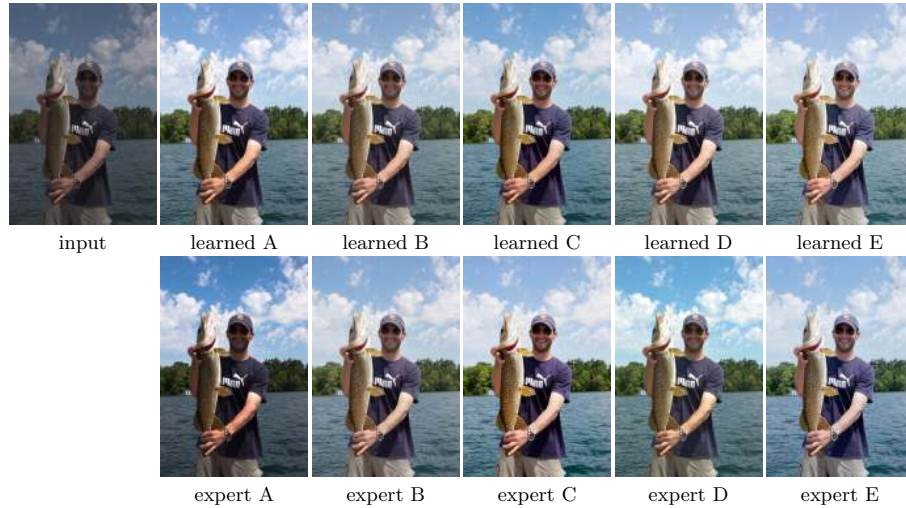


Fig. A3. Visual comparison of supervised learned mappings by our MT-GSGN model and expert labels.

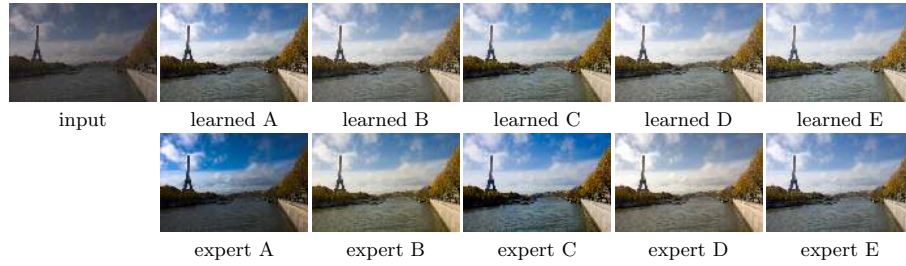


Fig. A4. Visual comparison of supervised learned mappings by our MT-GSGN model and expert labels.

A.3 Unsupervised learning - Visual comparison against the State-of-the-art

As the MIT5K [3] dataset has limited variation between the different experts in terms of style, we decided to train our MT-GSGN model on a custom made small flickr multi style dataset. Figures A6 to A13 show a visual comparison of learned unsupervised mappings by our MT-GSGN model against the state-of-the-art unsupervised methods DPE [19] and WhiteBox [10]. WhiteBox yields good results but often pushes the contrast too far. Our UL style result shows bright images but with limited contrast. The two HDR methods (DPE HDR and Ours HDR) both show visually pleasing results. While DPE has problems with artefacts in high contrast areas, especially on the border between the blue sky and the motiv, our method sometimes shows pixel level artefacts. We expect that these can be improved or resolved by fine tuning the training, as no extensive hyper-parameter search was performed in this example. Note that the two styles of the DPE method require two separate models and trainings, while our MT-GSGN method learns multiple styles in a single training using a compact model.

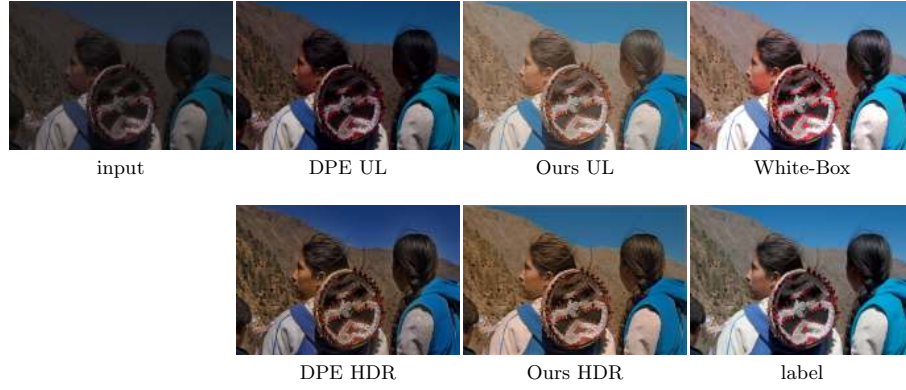


Fig. A5. Visual comparison of learned unsupervised mappings by our MT-GSGN model trained on the small Flickr Multi-style dataset, against DPE [19] and White-Box [10].

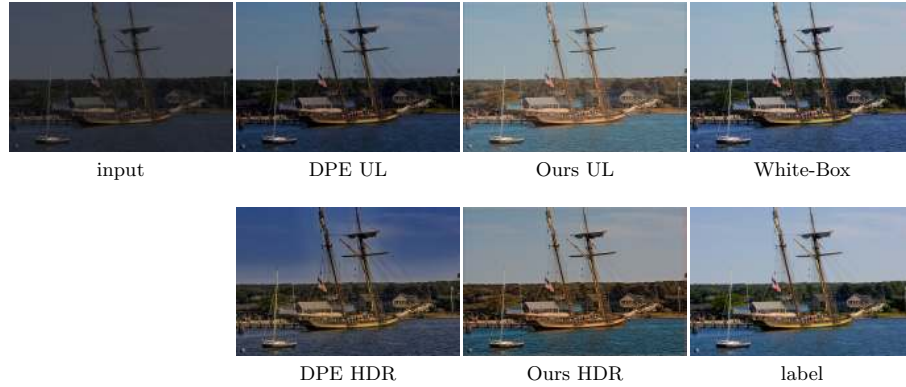


Fig. A6. Visual comparison of learned unsupervised mappings by our MT-GSGN model trained on the small Flickr Multi-style dataset, against DPE and WhiteBox.



Fig. A7. Visual comparison of learned unsupervised mappings by our MT-GSGN model trained on the small Flickr Multi-style dataset, against DPE and WhiteBox.

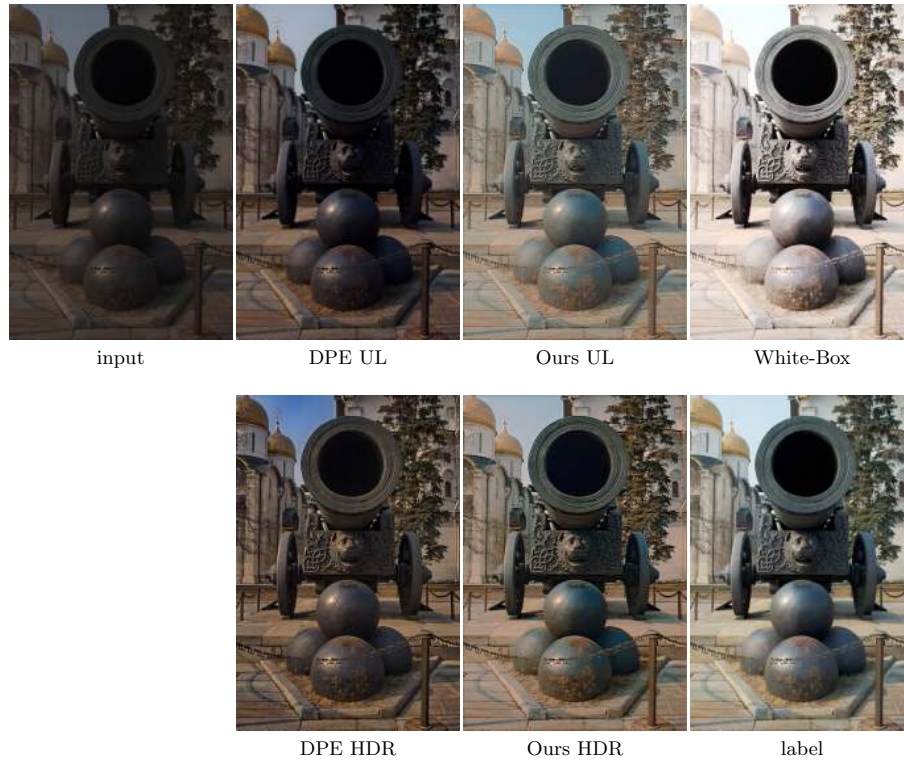


Fig. A8. Visual comparison of learned unsupervised mappings by our MT-GSGN model trained on the small Flickr Multi-style dataset, against DPE and WhiteBox.

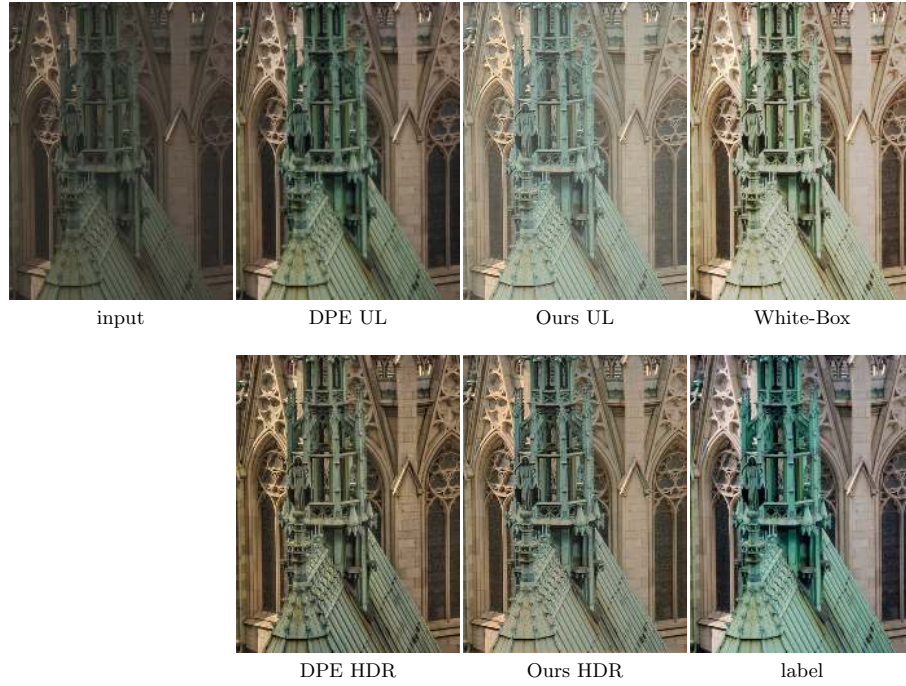


Fig. A9. Visual comparison of learned unsupervised mappings by our MT-GSGN model trained on the small Flickr Multi-style dataset, against DPE and WhiteBox.



Fig. A10. Visual comparison of learned unsupervised mappings by our MT-GSGN model trained on the small Flickr Multi-style dataset, against DPE and WhiteBox.



Fig. A11. Visual comparison of learned unsupervised mappings by our MT-GSGN model trained on the small Flickr Multi-style dataset, against DPE and WhiteBox.

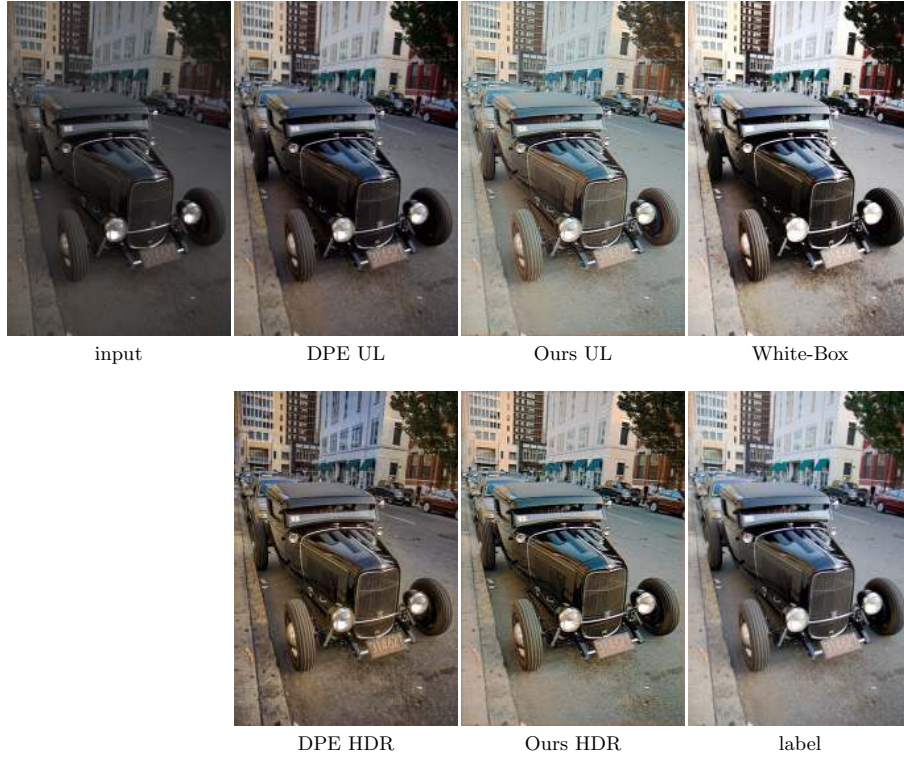


Fig. A12. Visual comparison of learned unsupervised mappings by our MT-GSGN model trained on the small Flickr Multi-style dataset, against DPE and WhiteBox.

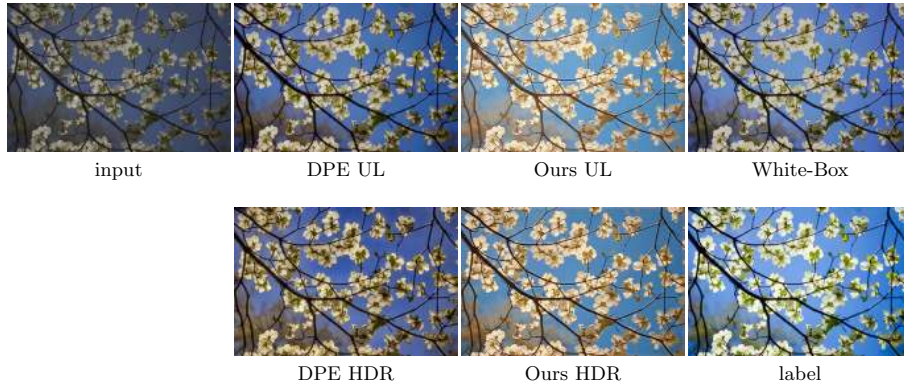


Fig. A13. Visual comparison of learned unsupervised mappings by our MT-GSGN model trained on the small Flickr Multi-style dataset, against DPE and WhiteBox.