

Color-wise Attention Network for Low-light Image Enhancement

Yousef Atoum
Yarmouk University

Mao Ye
BOSCH Research North America

Liu Ren
BOSCH Research North America

Ying Tai
Tencent YouTu

Xiaoming Liu
Michigan State University

Abstract

Absence of nearby light sources while capturing an image will degrade the visibility and quality of the captured image, making computer vision tasks difficult. In this paper, a color-wise attention network (CWAN) is proposed for low-light image enhancement based on convolutional neural networks. Motivated by the human visual system when looking at dark images, CWAN learns an end-to-end mapping between low-light and enhanced images while searching for any useful color cues in the low-light image to aid in the color enhancement process. Once these regions are identified, CWAN attention will be mainly focused to synthesize these local regions, as well as the global image. Both quantitative and qualitative experiments on challenging datasets demonstrate the advantages of our method in comparison with state-of-the-art methods.

1. Introduction

High-quality images carry rich information of the captured scene, enabling high-level computer vision tasks, such as object detection, object recognition, and scene understanding. However, challenges are often introduced in real-world environments making RGB-based perception difficult for both computers and humans. Among these challenges is the low-light condition in images captured in dark environments, such as night time or dark room images. They suffer from degraded brightness and contrast, added noise artifacts, and have a very narrow range of colors, such that understanding the original colors in the scene is a tedious task. E.g., in the low-light image of Fig. 1, many details such as the color of the building and trees are degraded and lost with the dark background.

Various methods have been proposed for the low-light image enhancement (LLIE), where researchers normally adopt a decomposition strategy for the sole purpose of sim-

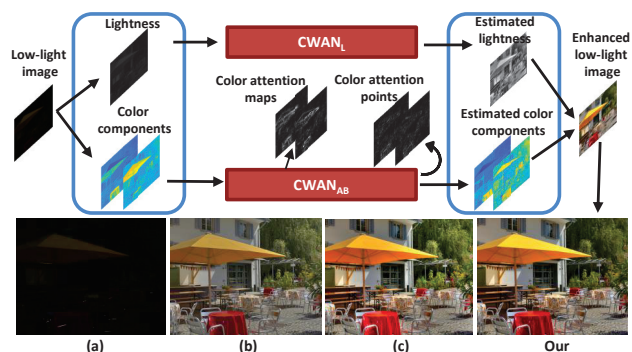


Figure 1. Our proposed CWAN method is illustrated in the top figure, where $CWAN_L$ enhances the lightness component and $CWAN_{AB}$ enhances the color components using an attention mechanism. We compare our result with (a) the low-light input, (b) LightenNet [19], and (c) ground truth.

plifying the problem. The decomposition strategy can be categorized into algorithm decomposition and image decomposition. Algorithm decomposition includes breaking the LLIE process into separate stages. For example, applying a denoising method prior to a lighting enhancement model [32], or vice-versa [13, 29]. Many Retinex-based methods attempt to estimate the illumination component first, and then the reflectance [22]. In [2], the enhancement is decomposed into three stages: luminance enhancement, detail enhancement, and a final enhancement stage. On the other hand, image decomposition includes breaking the image into multiple images where each holds a unique feature. For example, in [41, 1] the original low-light image along with the decomposed inverted image are used in a dehazing algorithm. In [35], they attempt to synthesize multi-exposure images to be used in a fusion algorithm.

In this paper, we propose to decompose the image into lightness and color components using the CIE LAB space, where each component is enhanced independently as seen in Fig. 1. Therefore, we use both algorithm and image decomposition strategies. Our main motivation is to simplify

the LLIE problem, by breaking it into two smaller problems. One is responsible for estimating the optimal lighting condition along with denoising given the dark image, while the other is required to revive the color information into its original state. Another motivation is the need to pay more attention to the color components. E.g., the estimated images from prior LLIE methods usually have degraded colors as in Fig. 1 (b).

The attention mechanism has been studied in a wide variety of computer vision problems, including object detection [37], tracking [16], segmentation [5], and action recognition [28]. These methods mimic the human visual system in making substantial use of contextual information in understanding RGB images. This involves discarding unwanted regions in the image, while focusing on more important parts containing rich features of our vision task. In this paper, we propose a novel color-wise attention CNN model, driven by key color features embedded in the low-light image. We hypothesize that these key colors can provide useful cues for image enhancement. These cues will be used as prior information in guiding and spanning the network’s attention to faithfully recover the color of the original image.

After decomposing a low-light image into lightness and color components, each component is enhanced independently via two CNN models named $CWAN_L$ and $CWAN_{AB}$ as in Fig. 1. Our $CWAN_{AB}$ computes color frequencies [17] in the dark image, and selects the colors with desired frequencies as the target. We experimentally show that, by ignoring high and low frequencies, we select colors belonging to regions of interest in an image, referred to as foreground colors in the scene. As humans, these foreground colors are what catches our eyes when first looking at dark images compared to background colors which has high frequency count. We learn $CWAN_{AB}$ to focus its attention in enhancing points belonging to these foreground colors. Note that the attention mechanism does not apply to $CWAN_L$, which performs lightness enhancement and denoising independently using a memory network structure [31]. We demonstrate state-of-the-art (SOTA) performance on real [3, 9] and synthetic [11] low-light datasets.

In summary, our main contributions are the following:

- ◊ Propose a novel color-wise attention network (CWAN) for LLIE. CWAN enhances the lightness image separately from enhancing the color component. By doing so we simplify the LLIE problem and achieve the state of the art.
- ◊ Propose a supervised attention mechanism utilizing color frequency maps in training $CWAN_{AB}$. With the color frequencies in an image, we select key local color points in the dark image which we desire CWAN to emphasize the enhancement on. Learning these selected colors is by nature a good starting point to spark the network’s attention.
- ◊ Achieve SOTA performance on several databases including both real and synthetic low-light images.



Figure 2. Decomposing an image into lightness and color components. (a) Input low-light image. (b) Lightness enhancement *only* using $CWAN_L$. (c) Color enhancement *only* using $CWAN_{AB}$. (d) The enhanced image from the proposed CWAN model.

2. Related Work

Generic Low-light enhancement methods Low-light image enhancement has been addressed over the past decades with various handcrafted techniques. A classical approach is to apply histogram equalization (HE), gamma correction, and their variations [23, 15]. Other researchers attempt more complex and global processing pipelines, e.g., estimating an illumination map via a Retinex-based method. LIME [13] and JED [26] both propose Retinex-based approaches for simultaneous LLIE and noise removal. Li et al. [20] additionally estimate a noise map in the Retinex pipeline. AMSR [18] proposes an adaptive multi-scale Retinex such that it assigns a weight to each single-scale Retinex output based on the image content.

Researchers adopt techniques from other low-level image enhancing disciplines. For example, Ying et al. [36] use the response characteristics of cameras for LLIE. BIMEF [35] and LECARM [27] both use the camera response model to synthesize multi-exposure images for fusion. By observing that the inverted low-light images intuitively look like haze images, the methods of [8, 41] apply image dehazing on the inverted image to enhance the image. All generic methods are based on handcrafted features and certain statistical models with many hyper-parameters. Thus, it is difficult for these methods to work in diverse real-world scenes.

Low-light enhancement via CNNs Deep neural network has witnessed success in low-level vision problems, e.g., super-resolution [6, 30], denoising [4], colorization [39, 38], dehazing [1, 25], multi-exposure fusion [24], and tone mapping [10]. CNNs have also been developed for low-light enhancement. LLNet [21] proposes a stacked sparse auto-encoder, for joint LLIE and denoising, which uses gray-scale images only without considering colors. Tao et al. [32] also propose a joint enhancement and denoising framework based on atmosphere scattering models. Cai et al. [2] utilize multi-exposure fusion (MEF) datasets to learn a joint luminance and detail enhancement CNNs. Similarly, MBLLN [12] uses multiple subnets and produces the enhanced image via multi-branch fusion. Wang et al. [34] first estimates an image-to-illumination mapping for modeling varying-lighting conditions and then takes this map to enhance underexposed images. SID [3] converts raw short-

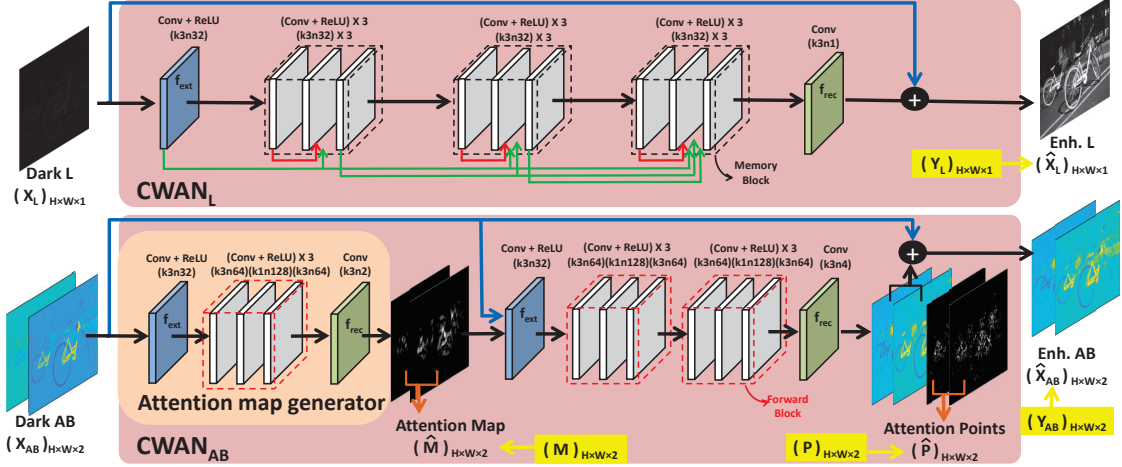


Figure 3. Network structure of the proposed CWAN method. Here, k3n64 indicates a kernel size of 3×3 , and a feature map number of 64. The stride is always equal to one in all layers. Red, green and blue arrows mean short, long and global skip connection representing local short-term memory, long-term memory, and residual learning respectively. The yellow arrow represents supervision in training CWAN.

exposure sensor images to RGB via a U-net-based denoising method followed by histogram stretching. Many recent methods utilize the Retinex theory to design the CNN model. E.g., LightenNet [19] and MSR-net [29] are based on the single and multi-scale Retinex models respectively, to estimate the illumination map via CNNs.

To the best of our knowledge, none of the prior CNN methods leverage color features in dark images for LLIE. Most CNN methods can enhance the image lightness to a great extent, however enhancing lightness alone is insufficient to generate high-quality RGB images with natural colors. Our proposed method introduces color attention maps from low-light images to be used as prior information for recovering natural high-quality images.

3. Proposed Method

Our color-wise attention model decomposes the low-light RGB image, into lightness and color components via the LAB color space. The motivation is to simplify the challenging LLIE process, and allow the color information drive the attention of $CWAN_{AB}$, while $CWAN_L$ focuses on enhancing image lightness and denoising simultaneously. As in Fig. 2, both the lightness and color components can be enhanced separately, and their fusion results in the final enhanced image. Enhancing both of these components separately is, conceptually, easier compared to performing both tasks at the same time.

3.1. Problem formulation

Given a low-light image, its lightness image $\mathbf{X}_L \in R^{H \times W}$ is fed to $CWAN_L$, and its color component images $\mathbf{X}_{AB} \in R^{H \times W \times 2}$ are fed to $CWAN_{AB}$. $CWAN_L$ outputs the enhanced lightness image $\hat{\mathbf{X}}_L \in R^{H \times W}$. $CWAN_{AB}$ outputs the enhanced color images $\hat{\mathbf{X}}_{AB} \in R^{H \times W \times 2}$, along with two intermediate outputs, the color attention maps

$\hat{\mathbf{M}} \in R^{H \times W \times 2}$ and sparse attention points $\hat{\mathbf{P}} \in R^{H \times W \times 2}$. We aim to train $CWAN_L$, denoted by $\mathcal{F}_L(\mathbf{X}_L; \theta_L)$, to map from low-light gray-scale image to an enhanced lightness image with reduced noise. Similarly, we train $CWAN_{AB}$, $\mathcal{F}_{AB}(\mathbf{X}_{AB}; \theta_{AB})$, to map from dark colors in low-light images to enhanced colors, under the constraint of color-wise attention. Thus, we formulate the LLIE problem as:

$$\theta_L^* = \arg \min_{\theta_L} E_{\mathbf{X}_L, \mathbf{Y}_L, \mathcal{D}_L} [\mathcal{L}_L(\mathcal{F}_L), \mathbf{Y}_L] \quad (1)$$

$$\theta_{AB}^* = \arg \min_{\theta_{AB}} E_{\mathbf{X}_{AB}, \mathbf{Y}_{AB}, \mathbf{P}, \mathcal{D}_{AB}} [\mathcal{L}_{AB}(\mathcal{F}_{AB}), \mathbf{Y}_{AB}, \mathbf{P}], \quad (2)$$

where \mathcal{D} denotes the training dataset, \mathcal{L} denotes the loss function, $\mathbf{Y}_L \in R^{H \times W}$ and $\mathbf{Y}_{AB} \in R^{H \times W \times 2}$ are the ground truth lightness and color components respectively, and $\mathbf{P} \in R^{H \times W \times 2}$ is a sparse set of ground truth attention points that are used to guide $CWAN_{AB}$. As seen from Fig. 3, $CWAN_{AB}$ firstly estimates attention maps $\hat{\mathbf{M}}$ with the attention map generator $\mathcal{F}_M(\mathbf{X}_{AB}; \theta_M)$ to help the learning of $\hat{\mathbf{P}}$. Before optimizing Eqn. 2, we pretrain $\mathcal{F}_M(\mathbf{X}_{AB}; \theta_M)$ to generate $\hat{\mathbf{M}}$, which is formulated as:

$$\theta_M^* = \arg \min_{\theta_M} E_{\mathbf{X}_{AB}, \mathbf{M}, \mathcal{D}_{AB}} [\mathcal{L}_M(\mathcal{F}_M), \mathbf{M}], \quad (3)$$

where $\mathbf{M} \in R^{H \times W \times 2}$ is the ground truth attention map. As described in Sec. 3.3, the attention maps and attention points serve different purposes. The former helps to identify local regions associated with foreground colors, while the latter specifies sparse points within the local region.

3.2. Network architecture

CWAN utilizes two fully convolutional networks (FCN), such that both FCNs are composed of a feature extraction conv layer (f_{ext}), several convolutional blocks in the middle, and a final feature reconstruction conv layer (f_{rec}). The

detailed structure is in Fig. 3. We learn the residual using a global skip connection, rather than the direct mapping, to ease the training difficulty. All blocks have the same number of conv and ReLU layers. Here, we define two types of blocks, memory blocks used in CWAN_L , and forward blocks used in CWAN_{AB} , as explained below.

3.2.1 CWAN_L structure

CWAN_L is composed of a series of memory blocks. These blocks are adopted from the image restoration work in [31], which was successfully used for image denoising, super-resolution and JPEG deblocking. We refer readers to [31], for a more detailed explanation on memory blocks. Generally, memory blocks utilize local short skip connections within the block to represent short-term memory, as well as long skip connections sourcing from previous blocks to represent long-term memory as in Fig. 3. The short- and long-term memory help CWAN_L to realize minor and major lightness enhancements within and across memory blocks.

3.2.2 CWAN_{AB} structure

In CWAN_L , all conv layers have the same number and size of filters, making long and short skip connection possible. In contrast, CWAN_{AB} does not utilize short and long skip connections. Instead, in each block, the middle conv layer is a nonlinear activation with 1×1 filters. This technique was successfully used in super-resolution [7].

The CWAN_{AB} network is composed of two parts. The first part $\mathcal{F}_M(\mathbf{X}_{AB}; \theta_M)$ takes \mathbf{X}_{AB} as input to generate an attention map $\hat{\mathbf{M}}$; the second part takes both \mathbf{X}_{AB} and $\hat{\mathbf{M}}$ forming a four-channel input to enhance colors. The goal of the first part is to internally estimate $\hat{\mathbf{M}}$ with high activations at points of interest in an image, such that it guides the local regions during the enhancement of the second part. Since \mathbf{X}_{AB} has two color channels, the estimated $\hat{\mathbf{M}}$ is also a two-channel attention map, so that attention is for each channel at each spatial coordinate. To supervise the learning of \mathcal{F}_M , we propose to use color frequency images to generate ground truth attention maps \mathbf{M} , described in Sec. 3.3. Therefore, the second part of CWAN_{AB} learns the mapping from the stacked \mathbf{X}_{AB} and $\hat{\mathbf{M}}$, to an enhanced $\hat{\mathbf{X}}_{AB}$ along with sparse attention color points $\hat{\mathbf{P}}$. The ground truth attention points \mathbf{P} are generated by selecting a set of non-zero foreground color points from \mathbf{M} . Both \mathbf{M} and \mathbf{P} play major roles in our color-wise attention mechanism.

3.3. Attention maps and points

A frequency image characterizes the spatial distribution along with the frequency information [17]. Given an image $\mathbf{X} \in R^{H \times W \times 3}$, we can compute its *color frequency image* $\mathbf{F} \in R^{H \times W \times 1}$, where $\mathbf{F}(x, y)$ equals to the number of occurrences of the RGB color $\mathbf{X}(x, y)$ in image \mathbf{X} . Examples of color frequency images are in Fig. 4.




Figure 4. (a) input image, (b) color frequency image \mathbf{F} , (c) attention map \mathbf{M} , and (d) attention points \mathbf{P} . For \mathbf{M} and \mathbf{P} , we show the first channel only.

We propose to utilize the color frequency image in CWAN_{AB} as follows. First, we apply a threshold τ on \mathbf{F} to eliminate specific undesired frequencies. E.g., the white background in the top image of Fig. 4 has very high frequency in \mathbf{F} . Whereas some parts of \mathbf{X} have very low frequencies in \mathbf{F} , e.g., noisy pixels or the eyes of the toy. By segmenting \mathbf{F} via $\tau_l < \mathbf{F} < \tau_u$, we emphasize our attention to foreground colors and eliminate both dominating color frequencies, and minor noisy regions. This results in a binary mask of desired color frequencies, denoted by $\bar{\mathbf{F}}$:

$$\bar{\mathbf{F}}(x, y) = \begin{cases} 1, & \text{if } \tau_l < \mathbf{F}(x, y) < \tau_u \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

After computing $\bar{\mathbf{F}}$, we generate the ground truth color attention map $\mathbf{M}(:, :, i) = \mathbf{X}_{AB}(:, :, i) \odot \bar{\mathbf{F}}$, where \odot is a Hadamard product, and $i \in [1, 2]$. We apply linear normalization on \mathbf{M} such that it is in the range of $[0, 1]$. Usually \mathbf{M} contains foreground colors as seen in Fig. 4.

In [40], authors learn a deep network for gray-scale image colorization, where the user can interact and guide the colorization process via *manually* selecting colors of specific pixels. Differs to the manual selection, CWAN_{AB} *randomly* selects a finite set of non-zero foreground color points β from \mathbf{M} to guide the color-wise attention model. We define a binary mask \mathbf{B}_P assigned with 1's at the coordinates of all β points, such that $\sum \sum \mathbf{B}_P = \beta$. Then, we compute the ground truth attention points \mathbf{P} to represent a sparse subset of foreground color points, via $\mathbf{P}(:, :, i) = \mathbf{M}(:, :, i) \odot \mathbf{B}_P$, where $i \in [1, 2]$.

In our attention mechanism, \mathbf{M} supervised the learning of $\hat{\mathbf{M}}$, which inputs to the subsequent network and guides CWAN_{AB} to focus on enhancing foreground colors at *coarse local regions*. On the other hand, \mathbf{P} guides the color enhancement at a *pixel level*, identifying key foreground colors. Since \mathbf{M} has much more duplicated foreground colors than \mathbf{P} , supervising by \mathbf{P} can cover majority of colors, yet using minimal, not duplicated, constraints.

3.4. Objective function

To train CWAN_L , we use the L_1 loss function to treat synthesizing the enhanced lightness $\hat{\mathbf{X}}_L$ as a regression

problem. On the other hand, training CWAN_{AB} takes place in two stages. The first stage trains the color-wise attention map generator to predict $\hat{\mathbf{M}}$ only, via the L_1 loss:

$$\mathcal{L}_M(\mathcal{F}_M(\mathbf{X}_{AB}; \theta_M), \mathbf{M}) = \|\hat{\mathbf{M}} - \mathbf{M}\|_1. \quad (5)$$

The second stage learns CWAN_{AB} end-to-end, including fine tuning the pretrained attention map generator. We propose to use the following loss function:

$$\mathcal{L}_{AB}(\mathcal{F}_{AB}(\mathbf{X}_{AB}; \theta_{AB}), \mathbf{Y}_{AB}) = \mathcal{L}_{\mathcal{H}} + \alpha \mathcal{L}_{MSE}. \quad (6)$$

Here α is the weight, $\mathcal{L}_{\mathcal{H}}$ is the Huber loss applied to $\hat{\mathbf{X}}_{AB}$. The Huber loss has witnessed great success in the image colorization field [38, 40], due to the relative high color saturation effect, making it suitable for enhancing low-light images. Further, we choose the Huber loss also because it is a robust estimator, and can help to avoid the averaging problem. The \mathcal{L}_{MSE} is the mean square error loss applied to the estimated attention points $\hat{\mathbf{P}}$, as follows:

$$\mathcal{L}_{\mathcal{H}} = \begin{cases} \frac{1}{2}(\hat{\mathbf{X}}_{AB} - \mathbf{Y}_{AB})^2, & \text{if } |\hat{\mathbf{X}}_{AB} - \mathbf{Y}_{AB}| \leq \delta \\ \delta|\hat{\mathbf{X}}_{AB} - \mathbf{Y}_{AB}| - \frac{1}{2}\delta^2, & \text{otherwise} \end{cases} \quad (7)$$

$$\mathcal{L}_{MSE} = \frac{1}{\beta} \|(\hat{\mathbf{P}} - \mathbf{P}) \odot \mathbf{B}_P\|_2^2, \quad (8)$$

where δ is the parameter of the Huber loss. While CWAN_{AB} outputs colors at all locations in $\hat{\mathbf{P}}$, \mathcal{L}_{MSE} loss is computed only using reconstructed colors at randomly sampled color locations in \mathbf{B}_P .

4. Experiments

This section provides ablation, quantitative and qualitative results, on both real-world and synthetic low-light images. We use metrics including Structural Similarity (SSIM), Peak Signal-to-Noise Ratio (PSNR), Lightness-Order-Error (LOE), Colorfulness (C), and a case study.

Datasets We use the See-In-the-Dark (SID) database [3] to learn CWAN, containing 5,094 raw low-light images, and 424 RGB ground truth images, such that multiple low-light images correspond to the same ground truth image. The data is divided into two distinct subsets, one captured with a Sony camera (SID_{Sony}), and another with a Fuji camera (SID_{Fuji}). We follow the same protocol in [3] to divide data into training, validating and testing sets. Since most low-light literature uses RGB as input, we convert the low-light images from raw to RGB using a COTS converter (easy2convert.com) which supports raw images of both Sony and Fuji cameras. The generated RGB images contain less noise than the original SID dataset, however, the low-light condition remains severe.

We further generate a synthetic low-light dataset from PASCAL VOC [11], which contains realistic scenes of objects. With randomly selected 1,000 images, we synthesize

Test image	HE	LLNet [21]	CWAN_L
Bird-D	11.28 / 0.62	18.43 / 0.60	28.76 / 0.91
Bird-D+GN	9.25 / 0.09	19.73 / 0.56	24.21 / 0.79
Girl-D	18.27 / 0.80	22.45 / 0.80	24.90 / 0.83
Girl-D+GN	16.07 / 0.26	20.04 / 0.60	21.96 / 0.77
House-D	12.03 / 0.65	21.10 / 0.64	21.55 / 0.70
House-D+GN	10.55 / 0.33	20.25 / 0.56	19.98 / 0.62
Pepper-D	18.45 / 0.85	21.33 / 0.78	25.48 / 0.83
Pepper-D+GN	14.69 / 0.21	22.33 / 0.78	24.25 / 0.80
Town-D	17.55 / 0.79	22.47 / 0.81	26.62 / 0.87
Town-D+GN	14.85 / 0.25	20.00 / 0.60	23.11 / 0.79

Table 1. Comparing lightness estimation in PSNR/SSIM. Here, D refers to a dark image generated using [21], and GN refers to Gaussian noise with $\sigma = 18$.

low-light images following a Retinex-based approach [19], with a 85% decrease in the pixel intensity. We refer to this as PASCAL₁₀₀₀ dataset. Finally, we evaluate CWAN using low-light images collected by [9] named HDRDB, which does not have groundtruth. With 96 low-light images of natural scenes, HDRDB has low-exposure, Gaussian noise, 5% of pixels are saturated and contain no information, with an additional 60% decrease on the pixel intensity using [19].

Experimental parameters For CWAN_L , we use three memory blocks, i.e., total of 11 conv layers, with the weight decay of 0.05 and batch size of 16. The large weight decay as in [7], helps to improve generalization on unseen low-light images. The number of patches from SID_{Sony} is 37,300 and from SID_{Fuji} is 33,100, of size 64×64 , i.e., 50 patches per training image. For CWAN_{AB} , we use one forward block in the attention map generator, i.e., total of 5 conv layers, and two forward blocks for the color enhancement part, i.e., total of 8 conv layers, with the weight decay of 0.05 and batch size of 32. The number of patches in training CWAN_{AB} from SID_{Sony} is 186,500 and from SID_{Fuji} is 165,500, of size 32×32 , i.e., 100 patches per training image. The patch sizes (32 or 64) were experimentally determined. We use $\tau_l = 0.05N$ and $\tau_u = 0.5N$ during training on small patches, where N is the number of pixels in the image. Other parameters in CWAN_{AB} are $\beta = 20$, $\alpha = 1$ and $\delta = 0.5$. The learning rate for all CNNs is 10^{-4} , and trained for 200 epochs. We implement CWAN on NVIDIA GTX1080Ti GPU, with Matconvnet toolbox [33].

4.1. Ablation Study

Analysis of CWAN_L We compare the performance of CWAN_L with LLNet [21] which uses a stacked auto-encoder for joint low-light enhancement and denoising, trained using gray-scale images¹, i.e., lightness enhancement only. We fine-tune CWAN_L using the same train set, on only 1,600 patches of size 64×64 for two epochs. We follow the same low-light test set generation as in [21], which only consists of 5 images. The results are shown in Tab. 1, and visual results are in the supplementary file. CWAN_L on average is able to outperform HE and LLNet

¹Dataset URL: <http://decsai.ugr.es/cvg/dbimagenes/>

α	1.4	1.2	1.0	0.8	0.6	0.4	0.2	0.0
PSNR	26.46	27.88	28.56	28.35	27.98	27.48	27.32	26.31

Table 2. Loss function weight analysis on $CWAN_{AB}$.

on all five images, demonstrating the effectiveness of lightness enhancement and denoising in $CWAN$. Nevertheless, the huge PSNR/SSIM margin in Tab. 1 is mainly due to differences in the network structures and training datasets.

Loss function weights of $CWAN_{AB}$ To understand the benefit of the attention model, we ablate the weights in the $CWAN_{AB}$ loss function in Eqn. 6. We first train $CWAN_{AB}$ with equal weights, $\alpha = 1$. Starting with this model, we fine-tune the network for 5 epochs independently with different α between 1.4 down to 0. This experiment is conducted on the train/test sets of SID_{Sony} . As in Tab. 2, our model heavily relies on the attention mechanism in low-light image enhancement. When the weights of both terms are equal, we reach a top PSNR of 28.56. When $\alpha > 1$, i.e., the attention points have more impact than the Huber loss in training $CWAN_{AB}$, the PSNR drops significantly. On the other hand, when $\alpha = 0$, i.e., $CWAN_{AB}$ is trained without the attention mechanism, the PSNR drops to 26.31.

Memory vs. forward blocks We adopt the memory and forward block concepts from low-level vision literature [31, 7], and incorporate them in $CWAN_L$ and $CWAN_{AB}$ respectively. We ablate all possible combinations of using two types of blocks in our model, and the PSNRs are as follows given the order of ($CWAN_L/CWAN_{AB}$): 27.23 with (forward/forward), 27.75 with (forward/memory), 28.56 with (memory/forward) and 28.17 with (memory/memory). Both types of blocks perform well in our $CWAN$, with the (memory/forward) configuration of ($CWAN_L/CWAN_{AB}$) being the best. Our method makes it possible to incorporate future novel block designs into our network.

Analysis on decomposing image into L and AB The intuitive alternative to our decomposition technique, is to learn a direct mapping from the input low-light color image to an enhanced RGB image. We train two baseline FCN models, using RGB and LAB input data named FCN_{RGB} and FCN_{LAB} respectively. Note that FCN_{LAB} estimates a LAB image, which is then converted to RGB. Both networks have the exact same structure as $CWAN_L$, with the same hyper-parameters. Using the test set of SID_{Sony} , FCN_{RGB} and FCN_{LAB} achieve a PSNR of 25.71 and 22.58 respectively, whereas $CWAN$ can reach a PSNR of 28.56. This shows how effective our decomposition technique over the traditional feed-forward FCN approach.

Alternatives of $CWAN_{AB}$ We introduce alternative methods of $CWAN_{AB}$ for comparison using the SID_{Sony} dataset. (1) $CWAN_{AB}^C$ removes the attention map generator, and only inputs X_{AB} to $CWAN_{AB}$. (2) $CWAN_{AB}^M$ inputs the attention map M , computed by the color frequency image F without the attention map generator, along with X_{AB}

Architectures	$CWAN_{AB}^C$	$CWAN_{AB}^M$	$CWAN_{AB}^P$	$CWAN_{AB}$
Input of CNN	X_{AB}	X_{AB} & M	X_{AB} & P	X_{AB}
Estimates \hat{M}	\times	\times	\times	\checkmark
Estimates \hat{P}	\times	\checkmark	\checkmark	\checkmark
Loss function	Eqn. 7	Eqn. 6	Eqn. 6	Eqn. 6
PSNR	25.86	27.93	27.39	28.56
SSIM	0.859	0.890	0.887	0.909

Table 3. Compare $CWAN_{AB}$ with alternative architectures.




Figure 5. M vs. \hat{M} : (a) a toy example, (b-d) low light images. The four middle columns are normalized in the range of $[0, 1]$, and represent the two channels of both M and \hat{M} respectively. the highest 30 responses in \hat{M} are overlaid in \hat{X} in the right column.

to $CWAN_{AB}$. (3) $CWAN_{AB}^P$ inputs the attention points P along with X_{AB} . All three architectures are trained on SID_{Sony} , while using the same $CWAN_L$ for lightness enhancement. As in Tab. 3, $CWAN_{AB}^C$ performs poorly without the attention map, but achieves higher PSNR than FCN_{LAB} which proves the usefulness of image decomposition. $CWAN_{AB}^M$ achieves the closest result to $CWAN_{AB}$ with a significant margin of 0.63 in PSNR. This shows that when training $CWAN_{AB}$ end-to-end, the attention map generator is able to learn a more effective \hat{M} than M . Thus, despite we could directly compute M from X_{AB} , it is better to let the attention map generator produce a better map \hat{M} for subsequent enhancement. $CWAN_{AB}^P$ proves that having a finite set of points guiding the network attention can achieve high PSNR, but not higher than feeding M .

Attention map analysis The attention map generator aims to select pixels with desired color frequencies. As in Fig. 5, our attention map generator is trained with M to generate \hat{M} highlighting foreground colors. On the other hand, after training $CWAN_{AB}$ end-to-end, the conv layers of the attention map generator were fine-tuned to estimates \hat{M} with high responses at colors of interest in the scene; it also tends to produce higher responses near object edges where gradients exist, as in (a). Note \hat{M} has no high responses in the background, as it occupies a large percentage of image, e.g., the street in (c). On the other hand, image (b) has various colors in the scene and \hat{M} produces high responses throughout the map, but tends to give the ball higher attention.

Attention point analysis P is constructed by selecting β random points from M . We study the impact of randomness




Figure 6. Color attention mechanism through the selected β points from \mathbf{P} in guiding the neighboring pixels.

in point selection, by independently fine-tuning CWAN_{AB} three times, where each time we regenerate the β points and train for five epochs. By testing the three models on the test set, the standard deviation of the PSNR is merely 0.002. This desired low impact is due to the controlled randomness, where we only select points from \mathbf{M} , but not \mathbf{X}_{AB} .

How does the color attention mechanism work? The random points selected in \mathbf{P} , while training CWAN_{AB} via the loss in Eqn. 8, has a major impact in reconstructing the colors in $\hat{\mathbf{X}}_{AB}$. To demonstrate the effectiveness of this attention mechanism, we study the effect of enhancing the colors at the β points compared to the local region surrounding those points. While training CWAN_{AB} end-to-end for various epochs, we compute the PSNR of only the β selected points as in Fig. 6, i.e., excluding all other pixels in the image. We compare the result with the PSNR of the 8-connected neighboring pixels surrounding the β points, i.e., 3×3 local regions excluding the center β points. Fig. 6 illustrates how CWAN_{AB} learns to enhance the colors at β points promptly, and gradually pulls the neighboring pixels for color enhancement.

One potential concern resides in the possibility of ignoring some colors in an image given the small value of β . To study how CWAN_{AB} selects a diverse set of colors, we first divide the colorful AB space into 40 clusters each with a unique color. Then we select a set of 2,000 training patches \mathbf{X}_{AB} along with their ground truth \mathbf{Y}_{AB} . For every \mathbf{Y}_{AB} , we assign all pixels to the corresponding cluster via nearest neighboring, forming a color map \mathbf{C} . For every \mathbf{X}_{AB} , we allow CWAN_{AB} to select $\beta = 20$ points at locations defined by \mathbf{B}_P . Based on the selected points, we evaluate the color clusters that have been selected from the non-zero elements of $\mathbf{C} \odot \mathbf{B}_P$, and compare them with all available colors in \mathbf{C} . By applying cross validation on three separate sets, the average percentage of selecting all possible colors is $88.3\% \pm 1.3\%$. This means that $\sim 90\%$ colors are selected by CWAN_{AB} , and the low standard deviation resembles consistency among the different sets. This study shows one strength of our color attention mechanism. That is, despite using a very small number of constraints, i.e., $\beta = 20$ points in \mathcal{L}_{MSE} , our optimization has an impact on the majority (90%) of to-be-recovered colors.

4.2. Quantitative results

We evaluate our method quantitatively on SID_{Sony} , SID_{Fuji} , PASCAL_{1000} , and HDRDB , comparing with 12 methods, including both generic and CNN-based LLIE methods. For comparison, we use the published codes for the first 11 methods, where SID is reimplemented by training a U-net structure with RGB data from SID_{Sony} via a L_1 loss. Tab. 4 summarizes the results.

The PSNR and SSIM are commonly used in LLIE literature to test image similarity with ground truth. Our method achieves the best result on all datasets except for the SSIM of PASCAL_{1000} , where LIME has the highest value. The LOE metric is commonly used for real low-light images [35, 36, 13]. Smaller LOE value means better lightness order is preserved, i.e., the intensity order of each pixel with all other pixels is similar between low-light and enhanced images. In Tab. 4, CWAN has the smallest LOE among all methods, demonstrating how well CWAN enhances low-light images while preserving the lightness order. We further utilize the colorfulness metric [14], which estimates the quality of colors in an image. Here, colorfulness represents the intensity and assortment of colors in the image, defined as $\sqrt{\sigma_{C_1}^2 + \sigma_{C_2}^2} + 0.37\sqrt{\mu_{C_1}^2 + \mu_{C_2}^2}$, where μ and σ are the mean and standard deviation of $C_1 = R - G$ and $C_2 = 0.5(R + G) - B$. Tab. 4 concludes that CWAN can synthesize more colorful images than all other methods. This proves how effective the proposed color-wise attention technique in faithfully restoring the color of the scene.

Finally, we conduct a user study on 20 low-light images in HDRDB which includes several indoor and outdoor scenes. We enhance the images with 12 baselines and CWAN resulting in a total of 13 images per low-light image. 10 people were asked to rank the images from the best (rank 1) to the worst image (rank 13). We summarize the scores in Tab. 4. The user study results demonstrate that CWAN has the best average rank of 2.32, and that it can generate visually appealing images, with great color attributes. CWAN also has a small standard deviation indicating the high consensus among 10 people.

4.3. Qualitative results

We qualitatively compare CWAN with the same baselines in Fig. 7 and 8, zooming into colorful regions where most baselines struggle in recovering the natural color. All of the methods tend to improve the lightness and colors of the low-light images at different performance levels. HE usually produces over or under-enhanced regions in the image due to the increase in global contrast. Dong, BIMEF, LECARM and Li tend to always have a grayish shade overlaid on the image. JED and AMSR perform well on denoising, but struggle with enhancing low-light. LightenNet enhances the brightness and contrast very well, but tends to generate white shadows very frequently, producing an un-

Method (training set)	PSNR / SSIM			LOE	Colorfulness	AVG rank	LLIE type
	SID _{Sony}	SID _{Fuji}	PASCAL ₁₀₀₀				
HE	21.68/0.712	20.25/0.726	26.28/0.885	46.53	14.87	5.49 ± 4.19	Generic
Dong [8]	20.56/0.824	22.40/0.841	20.79/0.907	78.88	28.16	9.60 ± 3.55	Generic
BIMEF [35]	15.55/0.697	16.52/0.708	14.13/0.732	38.78	28.23	9.94 ± 1.79	Generic
Ying [36]	18.61/0.746	20.12/0.751	18.24/0.810	36.25	13.06	9.58 ± 1.98	Generic
JED [26]	15.05/0.712	15.96/0.716	13.49/0.685	53.37	13.75	8.21 ± 1.95	Generic
AMSR [18]	12.86/0.567	12.31/0.414	12.38/0.626	318.83	16.29	10.21 ± 2.12	Generic
LIME [13]	19.10/0.802	21.39/0.830	20.72/0.935	105.23	32.02	7.28 ± 4.12	Generic
Li [20]	15.17/0.714	15.96/0.711	13.49/0.690	51.95	19.46	9.60 ± 2.38	Generic
LECARM [27]	17.99/0.831	18.39/0.829	16.86/0.884	32.83	19.17	5.16 ± 1.84	Generic
LightenNet [19]	19.21/0.817	17.87/0.812	17.82/0.858	257.04	28.76	3.07 ± 2.05	CNN
MBLLEN [12]	14.99/0.687	13.53/0.621	17.36/0.778	39.03	16.80	7.18 ± 2.18	CNN
SID [3] (SID _{Sony})	27.42/0.877	26.71/0.880	24.91/0.872	35.40	26.56	3.37 ± 1.43	CNN
CWAN (SID _{Sony})	28.56/0.909	28.11/0.911	29.08/0.924	30.53	40.38	2.32 ± 1.73	CNN
CWAN (SID _{Fuji})	27.28/0.902	26.77/0.911	28.68/0.923	—	—	—	CNN

Table 4. Quantitative results on four datasets. Red/blue fonts indicate the best/second best results.




Figure 7. Qualitative results of various methods compared to CWAN on SID_{Sony} (top), SID_{Fuji} (middle) and PASCAL₁₀₀₀ (bottom).




Figure 8. Qualitative results of various methods compared to CWAN on HDRDB.

pleasant visualization. LIME and MBLLEN method produces over-enhanced results especially in regions with original bright colors. SID recovers from dark images remarkable well, but CWAN surpasses U-net in enhancing more vibrant colors closer to ground truth.

5. Conclusions

This paper proposes a color-wise attention model for low-light image enhancement. The method attempts to

mimic the human visual system by first finding key colors in the dark image, and then spanning their attention spatially to generate a well enhanced image. The selected colors in the dark image are obtained by utilizing the color frequency map, and are by nature a good starting point to span the network's attention. The experimental results reveal the advantages of our method compared to SOTA low-light enhancement methods. Our method can produce visually pleasing and more realistic colors similar to the ground truth images.

References

- [1] B. Cai, X. Xu, K. Jia, C. Qing, and D. Tao. Dehazenet: An end-to-end system for single image haze removal. *IEEE Transactions on Image Processing*, 25(11):5187–5198, 2016. 1, 2
- [2] J. Cai, S. Gu, and L. Zhang. Learning a deep single image contrast enhancer from multi-exposure images. *IEEE Transactions on Image Processing*, 27(4):2049–2062, 2018. 1, 2
- [3] C. Chen, Q. Chen, J. Xu, and V. Koltun. Learning to see in the dark. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2, 5, 8
- [4] J. Chen, J. Chen, H. Chao, and M. Yang. Image blind denoising with generative adversarial network based noise modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3155–3164, 2018. 2
- [5] L. Chen, Y. Yang, J. Wang, W. Xu, and A. L. Yuille. Attention to scale: Scale-aware semantic image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3640–3649, 2016. 2
- [6] Y. Chen, Y. Tai, X. Liu, C. Shen, and J. Yang. FSRNet: End-to-end learning face super-resolution with facial priors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2492–2501, 2018. 2
- [7] C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):295–307, 2016. 4, 5, 6
- [8] X. Dong, G. Wang, Y. Pang, W. Li, J. Wen, W. Meng, and Y. Lu. Fast efficient algorithm for enhancement of low lighting video. In *2011 IEEE International Conference on Multimedia and Expo*, pages 1–6, July 2011. 2, 8
- [9] G. Eilertsen, J. Kronander, G. Denes, R. Mantiuk, and J. Unger. HDR image reconstruction from a single exposure using deep CNNs. *ACM Transactions on Graphics*, 36(6), 2017. 2, 5
- [10] Y. Endo, Y. Kanamori, and J. Mitani. Deep reverse tone mapping. *ACM Transactions on Graphics*, 36(6), 2017. 2
- [11] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. 2, 5
- [12] J. W. Feifan Lv, Feng Lu and C. Lim. MBLLEN: Low-light image/video enhancement using CNNs. *British Machine Vision Conference*, 2018. 2, 8
- [13] X. Guo, Y. Li, and H. Ling. LIME: Low-light image enhancement via illumination map estimation. *IEEE Transactions on Image Processing*, 26(2):982–993, 2017. 1, 2, 7, 8
- [14] D. Hasler and S. E. Suesstrunk. Measuring colorfulness in natural images. In *Human vision and electronic imaging VIII*, volume 5007, pages 87–96. International Society for Optics and Photonics, 2003. 7
- [15] H. Ibrahim and N. S. P. Kong. Brightness preserving dynamic histogram equalization for image contrast enhancement. *IEEE Transactions on Consumer Electronics*, 53(4), 2007. 2
- [16] S. E. Kahou, V. Michalski, and R. Memisevic. RATM: recurrent attentive tracking model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1613–1622, 2015. 2
- [17] T. Kashiwagi and S. Oe. Introduction of frequency image and applications. In *SICE, 2007 Annual Conference*, pages 584–591. IEEE, 2007. 2, 4
- [18] C.-H. Lee, J.-L. Shih, C.-C. Lien, and C.-C. Han. Adaptive multiscale retinex for image contrast enhancement. In *International Conference on Signal-Image Technology & Internet-Based Systems*, pages 43–50. IEEE, 2013. 2, 8
- [19] C. Li, J. Guo, F. Porikli, and Y. Pang. LightenNet: A convolutional neural network for weakly illuminated image enhancement. *Pattern Recognition Letters*, 104:15–22, 2018. 1, 3, 5, 8
- [20] M. Li, J. Liu, W. Yang, X. Sun, and Z. Guo. Structure-revealing low-light image enhancement via robust retinex model. *IEEE Transactions on Image Processing*, 27(6):2828–2841, 2018. 2, 8
- [21] K. G. Lore, A. Akintayo, and S. Sarkar. LLNet: A deep autoencoder approach to natural low-light image enhancement. *Pattern Recognition*, 61:650–662, 2017. 2, 5
- [22] S. Park, S. Yu, M. Kim, K. Park, and J. Paik. Dual autoencoder network for retinex-based low-light image enhancement. *IEEE Access*, 2018. 1
- [23] S. M. Pizer, E. P. Amburn, J. D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. ter Haar Romeny, J. B. Zimmerman, and K. Zuiderveld. Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*, 39(3):355–368, 1987. 2
- [24] K. R. Prabhakar, V. S. Srikanth, and R. V. Babu. Deepfuse: A deep unsupervised approach for exposure fusion with extreme exposure image pairs. In *IEEE International Conference on Computer Vision*, pages 4724–4732, 2017. 2
- [25] Y. Qu, Y. Chen, J. Huang, and Y. Xie. Enhanced pix2pix dehazing network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2019. 2
- [26] X. Ren, M. Li, W.-H. Cheng, and J. Liu. Joint enhancement and denoising method via sequential decomposition. In *IEEE International Symposium on Circuits and Systems*, pages 1–5. IEEE, 2018. 2, 8
- [27] Y. Ren, Z. Ying, T. H. Li, and G. Li. Lecarm: Low-light image enhancement using the camera response model. *IEEE Transactions on Circuits and Systems for Video Technology*, 29(4):968–981, 2018. 2, 8
- [28] S. Sharma, R. Kiros, and R. Salakhutdinov. Action recognition using visual attention. In *NIPS Time Series Workshop*. 2015. 2
- [29] L. Shen, Z. Yue, F. Feng, Q. Chen, S. Liu, and J. Ma. Msrnet: Low-light image enhancement using deep convolutional network. *arXiv preprint arXiv:1711.02488*, 2017. 1, 3
- [30] Y. Tai, J. Yang, and X. Liu. Image super-resolution via deep recursive residual network. In *Proceeding of IEEE Computer Vision and Pattern Recognition*, Honolulu, HI, July 2017. 2
- [31] Y. Tai, J. Yang, X. Liu, and C. Xu. Memnet: A persistent memory network for image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4539–4547, 2017. 2, 4, 6

- [32] L. Tao, C. Zhu, J. Song, T. Lu, H. Jia, and X. Xie. Low-light image enhancement using cnn and bright channel prior. In *IEEE International Conference on Image Processing*, pages 3215–3219. IEEE, 2017. [1](#), [2](#)
- [33] A. Vedaldi and K. Lenc. Matconvnet: Convolutional neural networks for MATLAB. In *Proceedings of the 23rd ACM International Conference on Multimedia*, pages 689–692. ACM, 2015. [5](#)
- [34] R. Wang, Q. Zhang, C.-W. Fu, X. Shen, W.-S. Zheng, and J. Jia. Underexposed photo enhancement using deep illumination estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2019. [2](#)
- [35] Z. Ying, G. Li, and W. Gao. A bio-inspired multi-exposure fusion framework for low-light image enhancement. *arXiv preprint arXiv:1711.00591*, 2017. [1](#), [2](#), [7](#), [8](#)
- [36] Z. Ying, G. Li, Y. Ren, R. Wang, and W. Wang. A new low-light image enhancement algorithm using camera response model. In *IEEE International Conference on Computer Vision Workshops*, pages 3015–3022, Oct 2017. [2](#), [7](#), [8](#)
- [37] D. Yoo, S. Park, J.-Y. Lee, A. S. Paek, and I. So Kweon. Attentionnet: Aggregating weak directions for accurate object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2659–2667, 2015. [2](#)
- [38] S. Yoo, H. Bahng, S. Chung, J. Lee, J. Chang, and J. Choo. Coloring with limited data: Few-shot colorization via memory augmented networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2019. [2](#), [5](#)
- [39] R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. In *European Conference on Computer Vision*, pages 649–666. Springer, 2016. [2](#)
- [40] R. Zhang, J.-Y. Zhu, P. Isola, X. Geng, A. S. Lin, T. Yu, and A. A. Efros. Real-time user-guided image colorization with learned deep priors. *ACM Transactions on Graphics*, 9(4), 2017. [4](#), [5](#)
- [41] X. Zhang, P. Shen, L. Luo, L. Zhang, and J. Song. Enhancement and noise reduction of very low light level images. In *21st International Conference on Pattern Recognition*, pages 2034–2037. Ieee, 2012. [1](#), [2](#)