

Better Compression with Deep Pre-Editing

Hossein Talebi, Damien Kelly, Xiyang Luo, Ignacio Garcia Dorado,
Feng Yang, Peyman Milanfar and Michael Elad
Google Research - Mountain View, California

Abstract

Could we compress images via standard codecs while avoiding artifacts? The answer is obvious – this is doable as long as the bit budget is generous enough. What if the allocated bit-rate for compression is insufficient? Then unfortunately, artifacts are a fact of life. Many attempts were made over the years to fight this phenomenon, with various degrees of success. In this work we aim to break the unholy connection between bit-rate and image quality, and propose a way to circumvent compression artifacts by pre-editing the incoming image and modifying its content to fit the given bits. We design this editing operation as a learned convolutional neural network, and formulate an optimization problem for its training. Our loss takes into account a proximity between the original image and the edited one, a bit-budget penalty over the proposed image, and a no-reference image quality measure for forcing the outcome to be visually pleasing. The proposed approach is demonstrated on the popular JPEG compression, showing savings in bits and/or improvements in visual quality, obtained with intricate editing effects.

1. Introduction

Commonly used still image compression algorithms, such as JPEG [52], JPEG-2000 [12], HEIF [1] and WebP [22] produce undesired artifacts when the allocated bit rate is relatively low. Blockiness, ringing, and other forms of distortion are often seen in compressed-decompressed images, even at intermediate bit-rates. As such, the output images from such a compression procedure are of poor quality, which may hinder their use in some applications, or more commonly, simply introduce annoying visual flaws.

Numerous methods have been developed over the years to confront this problem. In Section 4 we provide a brief review of the relevant literature, encompassing the various strategies taken to fight compression artifacts. Most of the existing solutions consider a post-processing stage that removes such artifacts after decompression [3, 4, 11, 13, 17, 18, 30, 31, 34, 42, 48, 54, 57–59]. Indeed, hundreds of papers

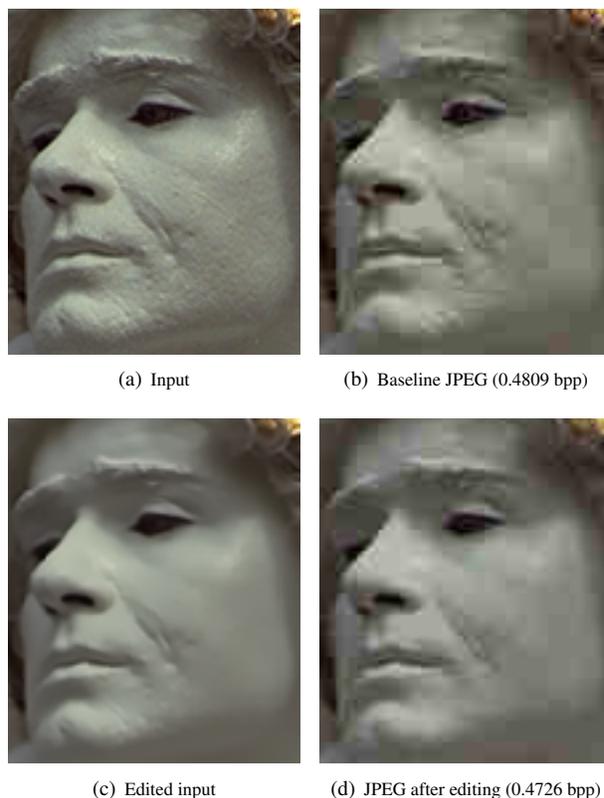


Figure 1: Comparison of our pre-editing method with baseline JPEG. The uncompressed input (a) is compressed by JPEG (b), which shows a lot of compression artifacts. We propose to edit the input image (c) before JPEG compression (d) to obtain a better perceptual quality and lower bit rate.

that take this post-processing approach have been published over the years, including recent deep-learning based solutions (e.g., [10, 16, 20, 53]).

Far less popular are algorithms that propose to pre-process the image prior to its compression, in order to reduce its entropy, thus avoiding the creation of artifacts in the first place [14, 15, 33, 40, 41, 50]. Indeed, a denoising applied before the compression is often found effective for better encoding performance (e.g. [45]). This line of thinking is

scarcer in the literature due to the more complex treatment it induces and the weaker control it provides on the output artifacts. Still, such a pre-processing approach has a great advantage over the alternatives, as the changes to the image are done on the server side, while the decoder side does not need to be modified nor adjusted.

In this work we propose to pre-process the image by automatically editing its content, applied before its compression using a standard coding algorithm. Our goal is to modify the image content smartly so as to guarantee that (i) most of the visual information in the image is preserved; (ii) the subsequent compression operates in a much better regime and thus leads to reduced artifacts; and (iii) the edited image after compression is still visually appealing. By considering all these forces holistically, we aim to get creative editing effects that enable the compression-decompression stage to perform at its best for the given bit budget.

While one could pose the proposed editing task as an optimization problem to be solved for each incoming image separately, we take a more challenging route, in which we target the design of a universal deep neural network that performs the required editing on any input image. The clear advantage in this approach is the speed with which inference is obtained once the network has been trained.

Our learning relies on minimizing a loss-function that includes three key penalties, aligned with the above description. The first forces the original and the edited images to be “sufficiently close” to each other, while still allowing content editing. A second term penalizes the bit content of the edited image, so as to force the bit-budget constraint while striving for an artifact-free compression. This part is achieved by yet another network [5] that predicts the entropy and quality of the image to be compressed. Last, but definitely not least, is a third penalty that encourages the edited image after compression to be visually pleasing. Indeed, our formulation of the problem relies heavily on the availability of a no-reference quality metric, a topic that has seen much progress in recent years [8,27,37–39,47,56]. All the above-mentioned ingredients are posed as differentiable machines, enabling an end-to-end effective learning of the editing operation. An example of the proposed technique is shown in Fig. 1, where the editing operation allows for better perceptual quality and lower bit budget.

2. Formulating the Problem

We start with a few definitions that will help in formulating our problem and its solution.

Definition 1 (Codec Operation) We define by $C_R(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}^N$ the process of compression and decompression of a given image with R bits. This function gets an image \mathbf{x} and produces an image, $C_R(\mathbf{x})$, possibly with the compression artifacts mentioned above.

Definition 2 (Quality Assessment) We define by $Q(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}^+$ the process of allocating a no-reference quality to a given image \mathbf{x} . The output is a non-negative scalar with values tending to zero for higher quality images.

Definition 3 (Distance Measure) We define by $d(\mathbf{x}_1, \mathbf{x}_2)$ the distance between two images, \mathbf{x}_1 and \mathbf{x}_2 , of the same size. Our distance function should be defined such that it is “forgiving” to minor content changes such as small geometrical shifts or warps, delicate variations in gray-values, or removal of fine texture.

Armed with the above, we are now ready to formulate our problem. Given an image \mathbf{z} to be compressed with a bit budget of R bits, the common practice is to perform compression and decompression directly, $\hat{\mathbf{z}} = C_R(\mathbf{z})$, and live with the limitations.

In this work we suggest a novel alternative: We seek a new image \mathbf{x} that is (i) as close as possible to the given image \mathbf{z} ; (ii) it is compressible using R bits; and most importantly (iii) it is of high quality. Naturally, \mathbf{x} will be an edited variation of \mathbf{z} in which some of the content has been changed, so as to enable good quality despite the compression. Here is our first attempt to formulate this problem:

$$\min_{\mathbf{x}} \text{dist}(\mathbf{x}, \mathbf{z}) + \lambda Q(\mathbf{x}) \quad \text{s.t.} \quad \mathbf{x} = C_R(\mathbf{x}). \quad (1)$$

In words, given \mathbf{z} and R we seek an image \mathbf{x} that is close to \mathbf{z} , it is of high quality (low value of $Q(\mathbf{x})$), and it can be represented via R bits. Referring to the constraint, recall that the compression-decompression operation is idempotent, i.e. applying it more than once on a given image results with the same outcome as using it once [29]. Thus, the constraint aims to say that \mathbf{x} is a feasible outcome of the compression algorithm with the given budget of R bits.

An alternative formulation that may serve the same goal is one in which we fix the quality as a constraint as well,

$$\min_{\mathbf{x}} \text{dist}(\mathbf{x}, \mathbf{z}) \quad \text{s.t.} \quad \mathbf{x} = C_R(\mathbf{x}) \text{ and } Q(\mathbf{x}) = Q_0, \quad (2)$$

so as to say that whatever happens, we insist on a specific output quality, willing to sacrifice content accordingly.

Both problems defined in Equations (1) and (2), while clearly communicating our goal, are hard to handle. This is mainly due to the non-differentiable nature of the function $C_R(\mathbf{x})$, and the fact that it is hard to fix a rate R while modifying the image \mathbf{x} . While these could be dealt with by a projection point of view (see [6]), we take a different route and modify our formulation to alleviate these difficulties. This brings us to the following additional definitions:

Definition 4 (Quality-Driven Codec Operation) We define by $C_q(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}^N$ the process of compression and decompression of a given image with a quantization

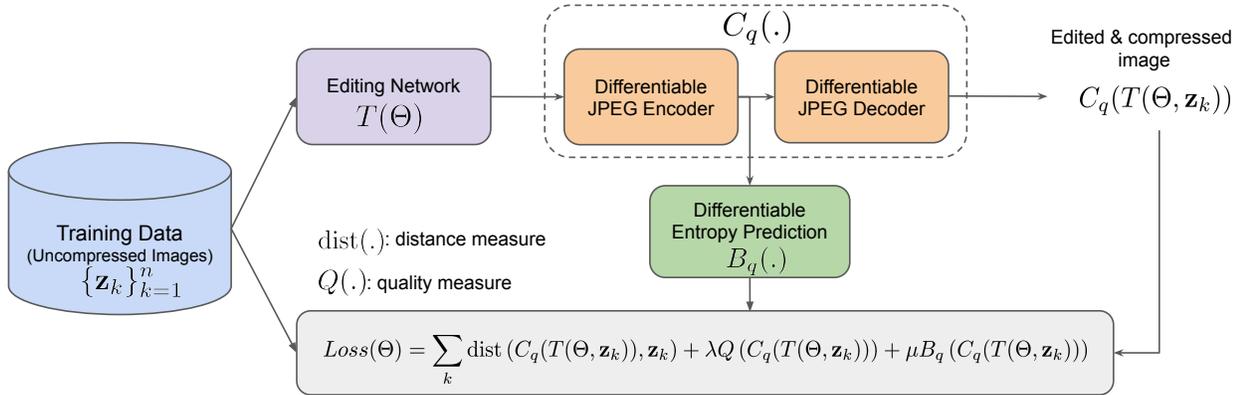


Figure 2: Our learning pipeline for training the image editing $T(\Theta)$. Input image \mathbf{z}_k is first edited by our editing network. Then, the edited image is fed to the differentiable JPEG encoder/decoder. The entropy of the quantized DCT coefficients are predicted and used in our training loss. To ensure that the compressed image is close to the uncompressed input, we use a distance measure. We also use a quality term to enforce the human perceptual preference.

(or quality factor) q . This function gets an image \mathbf{x} and produces an image, $C_q(\mathbf{x})$, possibly with the compression artifacts mentioned above.

Definition 5 (Entropy Predictor) We define by $B_q(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}^+$ the process of predicting the compression-decompression performance of a specific algorithm (e.g., JPEG) for a given image \mathbf{x} and a quantization level q . This function produces the expected file-size (or entropy).

Note that by setting q , we induce a roughly fixed PSNR on the image after compression-decompression. Thus, by minimizing $B_q(\mathbf{x})$ with respect to \mathbf{x} , we will be aiming for reducing the file size while preserving quality. Returning to our formulation, we add a penalty term, $B_q(\mathbf{x})$, so as to guarantee that the rate is preserved (or more accurately, controlled). This leads to

$$\min_{\mathbf{x}} \text{dist}(\mathbf{x}, \mathbf{z}) + \lambda Q(\mathbf{x}) + \mu B_q(\mathbf{x}) \quad \text{s.t.} \quad \mathbf{x} = C_q(\mathbf{x}). \quad (3)$$

The constraint $\mathbf{x} = C_q(\mathbf{x})$ assures that \mathbf{x} is a valid output of the compression, and it can be alternatively written as¹

$$\min_{\mathbf{x}} \text{dist}(C_q(\mathbf{x}), \mathbf{z}) + \lambda Q(C_q(\mathbf{x})) + \mu B_q(C_q(\mathbf{x})). \quad (4)$$

If we have a differentiable proxy for the compression operation $C_q(\cdot)$, the above loss is manageable.

We could have stopped here, handling this optimization task and getting an edited and compressed image \mathbf{x} for any incoming image \mathbf{z} . This could have been a worthy and even fascinating feat by itself, but we leave it for future work.

¹Admittedly, the notations introduced are a bit cumbersome, as both B and C use the same quantization level q . The alternative could have been to divide the compression C_q into an encoder and decoder, and feed the encoder result to B without specifying q . We chose to stay with the above formulation for consistency with the opening description.

As we have already mentioned in the introduction, we aim higher. Our goal is to design a feed-forward CNN that would perform this editing for any given image automatically. Denoting this editing network by $T(\Theta, \mathbf{z})$, where Θ are the network parameters to be learned/set, our training loss is given by the following expression:

$$Loss(\Theta) = \sum_k [\text{dist}(C_q(T(\Theta, \mathbf{z}_k)), \mathbf{z}_k) + \lambda Q(C_q(T(\Theta, \mathbf{z}_k))) + \mu B_q(C_q(T(\Theta, \mathbf{z}_k)))] \quad (5)$$

This expression simply sums the per-image loss over many training images $\{\mathbf{z}_k\}$, and replaces the edited image \mathbf{x}_k by the network's output $T(\Theta, \mathbf{z}_k)$. Minimizing this loss with respect to Θ , we obtain the editing network, as desired. Our learning pipeline is shown in Fig. 2.

3. The Proposed Approach

In this section we dive into our implementation of the above-discussed editing idea. We start by specifying the ingredients used within the loss function and then turn to describe the training procedure employed.

3.1. Specifying the Training Loss

Returning to our definitions, we would like to identify the specific ingredients used in Equation (5). In this work we concentrate on the JPEG compression algorithm, due to its massive popularity and the fact that it is central in digital imaging and mobile photography. Our formulation relies on three ingredients:

Distance measure: We seek a definition of $\text{dist}(\mathbf{x}_1, \mathbf{x}_2)$ that does not penalize moderate editing changes between the two images. In our implementation we construct this distance measure by feature extraction function, $F(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}^L$, and use the perceptual loss $\|F(\mathbf{x}_1) - F(\mathbf{x}_2)\|_1$

as our distance [21, 28, 51]. These features could be the activations of an inner layer within the VGG-16 [44] or the NIMA [47] networks, and they are used to characterize an image in a domain that is less sensitive to the allowed perturbations. The deeper these features are taken from, the more daring the editing of the image is expected to be. We experimented with various VGG-16 activations trained for image quality assessment task [47], and selected the output of the second convolutional layer before max pooling as our feature extraction function $F(\mathbf{x})$.

Quality measure: We assess image quality using NIMA [47]. NIMA is a no-reference image quality assessment machine that has been used for training image enhancement [46].

Differentiable JPEG: As mentioned above, we need to incorporate the function $C_q(\cdot)$ within our loss and thus it should be differentiable. Indeed, as we are working with JPEG, this function does not control the rate but rather the quality factor q when running this compression-decompression. We obtain a differentiable version of this operator by replacing the quantization step-function curve by a 3^{rd} -degree smoothed polynomial. Our implementation in essence is quite similar to [43].

Entropy prediction: In our framework with JPEG, the discrete entropy of the quantized DCT coefficients should be measured. However, just as described above, the derivatives of the quantization operation are zero almost everywhere, and consequently gradient descent would be ineffective. To allow optimization via stochastic gradient descent, we use the entropy estimator proposed in [5], where an additive i.i.d. uniform noise is added to the quantized coefficients. This means that the probability mass function of the DCT coefficients is estimated by a continuous relaxation of it, implying that the differential entropy of the DCT coefficients can be used as an approximation of the discrete entropy. This provides a slightly biased estimate of the discrete entropy in coarser quantization regimes, but the bias vanishes for finer quantization levels [5].

To summarize, the following is the loss function we use in our experiments:

$$Loss(\Theta) = \sum_k [\text{dist}(C_q(T(\Theta, \mathbf{z}_k)), \mathbf{z}_k) + \lambda Q(C_q(T(\Theta, \mathbf{z}_k))) + \mu B_q(C_q(T(\Theta, \mathbf{z}_k)))]$$

where the distance function represents the perceptual error measures, the image quality is computed by NIMA and a total variation measure, and the entropy estimate is computed over the quantized DCT coefficients of the edited image. Note that the same q-factor is applied both in the function $B_q(\cdot)$ and the differentiable JPEG function $C_q(\cdot)$.

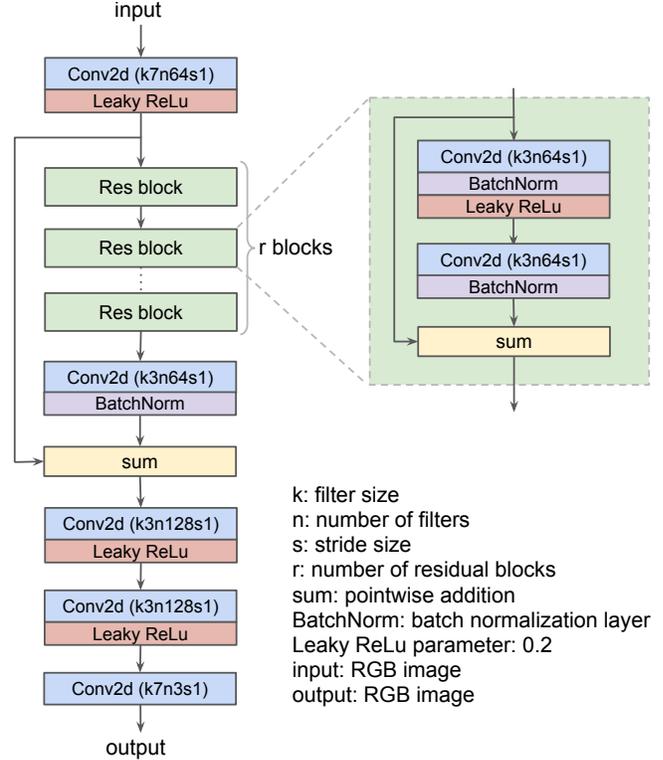


Figure 3: Our image smoothing CNN.

3.2. The Editing Network

Our editing network consists of two parts; An image smoothing (Fig. 3), and a patch-based warping operation (Fig. 4). While the smoothing is similar to a denoiser that controls fine-grained details, the spatial transformer allows for subtle local warps that make the underlying image more compressible [41]. More details on both parts is given below.

3.2.1 The Smoothing Network

Our image smoothing CNN is shown in Fig. 3. This convolutional neural network is similar to the residual CNN of Ledig et al. [35]. This architecture has $r = 2$ identical residual blocks, with 3×3 kernels and 64 feature maps followed by batch normalization [25] layers and Leaky ReLu (instead of parametric ones) [23] activations. To avoid boundary artifacts, the input image and feature maps are symmetrically padded before convolutions.

Examples of using the trained smoothing network are shown in Fig. 5. These images are not compressed by JPEG, and only represent edits applied to the input. The difference image shows that our editing removes fine details. Note that compressing the smoothed image with JPEG encoder at quality factor 20 takes 1.03 bpp, whereas the same

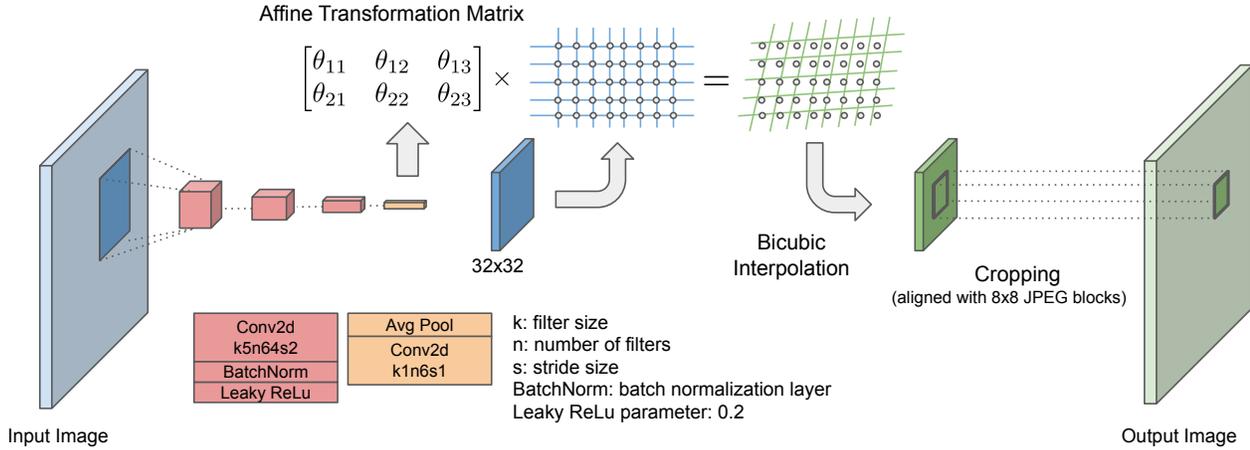


Figure 4: Our patch-based spatial transformer network. The affine transformer parameters of 32×32 blocks are obtained from a trainable CNN. Transformed image grid is interpolated to obtain a warped image block of size 32×32 . Finally, an 8×8 central block is extracted.

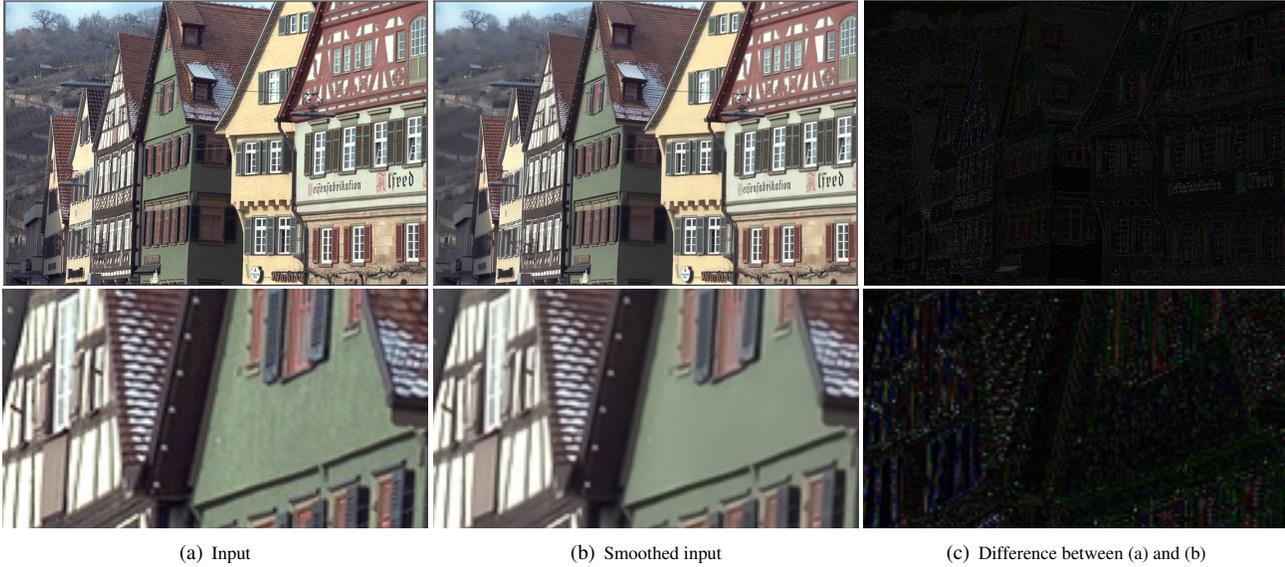


Figure 5: The difference between the input and the smoothed images (without JPEG compression). Our smoothing trained-network removes fine-grain details from the input image to make it more compressible by JPEG. Compressing (a) and (b) images with JPEG encoder at quality factor 20 takes 1.15 and 1.03 bpp, respectively.

encoder takes 1.15 bpp for compressing the input image.

3.2.2 The Spatial Transformer Network (STN)

As shown by Rott et al. [41], local image deformations can lead to DCT domain sparsity and consequently better compressibility. Unlike [41] that solves an alternating optimization with an optical flow, we use the differentiable spatial transformer network [26]. STN learns 6 parameters for an affine local transformation that allows cropping, translation, rotation, scale, and skew to be applied on the input (Fig. 4).

We apply STN on overlapping blocks of size 32×32 , and then we extract central crops of size 8×8 that are aligned with JPEG blocks. Since each 32×32 block is warped separately, this can cause inconsistency near the boundary of cropped blocks. To alleviate this, all overlapped grid values are averaged across neighboring blocks.

Examples of using the trained STN are shown in Fig. 6. The STN warps textures and edges locally to make the 8×8 blocks more compressible by JPEG encoder. Compressing the input and deformed images in Fig. 6(a) and Fig. 6(b) with JPEG encoder at quality factor 20 requires 0.725 bpp



(a) Input

(b) Warped input

(c) Difference of (a) and (b)

Figure 6: The difference between the input and the warped images (without JPEG compression). Our warping makes spatial transformations on local image patches to make them more compressible by JPEG. Compressing (a) and (b) images with JPEG encoder at quality factor 20 takes 0.725 and 0.708 bpp, respectively.

and 0.708 bpp, respectively.

To take advantage of both editing stages, we cascade the smoothing and warping operations. While the smoothing allows for less blockiness artifacts, the STN leads to texture preservation. Next, we discuss our training data.

3.3. Data

Our editing networks are trained on uncompressed images. To this end, we use burst processed images of Hasi-noff et al. [24], which provides 3640 images of size 12 mega pixels. All images are converted to 8-bit PNG format. We extract about 120K non-overlapping patches of size 480×640 , and use them to train our model. We also create a test set with 10% of the data.

4. Relation to Prior Work

We pause our main story for a while and discuss the rich literature on combating compression artifacts. Our goal is to give better context to the suggested methodology by presenting the main existing alternatives. Note, however, that this does not constitute an exhaustive scan of the existing literature, as this is beyond the scope of this work. We survey these algorithms by dividing them into categories based on their core strategies:

Post-Processing Algorithms [3, 4, 11, 13, 17, 18, 30, 31, 34, 42, 48, 54, 57–59]: Those are the most common methods available, operating on the image after the compression-decompression damage has already been induced. Algorithms of this sort that are designed in the context of the JPEG format are known as deblocking algorithms. The idea behind these methods, be it for JPEG or any

other transform-based coder, is quite simple, even though there are many ways to practice it; Given the compressed-decompressed image and knowing the quantization levels and the transform applied, the original image to be recovered must lie in a convex set that has a rotated hyper-rectangle shape. A recovery algorithm should seek for the most probable image within this set, something that could be done by relying on various regularization strategies. While some algorithms make use of this very rationale directly, others relax it in various ways, by simplifying the constraint set to a sphere, by forcing the recovery algorithm to take a specific shape, and more. At its simplest form, such a deblocking could be a simple linear filter applied to the boundaries between adjacent blocks.

Deep-Learning Based Solutions [10, 16, 20, 53]: Still under the regime of post-processing, recent solutions rely on deep neural networks, trained in a supervised fashion to achieve their cleaning goal. These methods tend to be better performing, as their design targets the recovery error directly, instead of relying on model-based restoration methods.

Scale-Down and Scale-Up [9, 36, 49, 55]: An entirely different way to avoid compression artifacts is to scale-down the image before compression, apply the compression-decompression on the resulting smaller image, and scale-up the outcome in the client after decompression. This approach is especially helpful in low bit-rates, since the number of blocks is reduced, the bit stream overhead reduce along with it, and the scale-up at the client brings an extra smoothing. Variations over this core scheme have been proposed over the years, in which the scale-down or up are optimized for better end-to-end performance.

Pre-Processing Algorithms [33, 40, 45, 50]: It is well known that compression-decompression often behaves as a denoiser, removing small and faint details from the image. Nevertheless, applying a well-designed denoiser prior to the compression may improve the overall encoding performance by better prioritizing the content to be treated. The existing publications offering this strategy have typically relied on this intuition, without an ability to systematically design the pre-filter for best end-to-end performance, as the formulation of this problem is quite challenging.

Deformation Aware Compression [41]: While this work offers a pre-processing of the image along the same lines as described above, we consider it as a class of its own because of two reasons: (i) Rather than using a denoiser, the pre-process applied in this work is a geometrical warp, which re-positions elements in the image to better match the coder transform and block-division; and (ii) the design of the warp is obtained by an end-to-end approximate optimization method. Indeed, this paper has been the source of inspiration behind our ideas in this work.

Our proposed method joins the list of pre-processing based artifact removal algorithms, generalizing the work in [41] in various important ways: (i) Our method could accommodate more general editing effects; (ii) Its application is simple and fast, once the editing network has been trained; and (iii) We employ a no-reference image quality assessment that supports better quality outcomes. As already mentioned, the pre-processing strategy has a unique advantage over the alternative methods in the fact that the decoder does not have to be aware of the manipulations that the image has gone through, applying a plain decoding, while leaving the burden of the computations to the encoder. That being said, we should add that this approach can be easily augmented with a post-processing stage, for fine-tuning and improving the results further.

We conclude this survey of the relevant literature by referring to two recent and very impressive papers. The work reported in [7] offers a theoretical extension of the classic rate-distortion theory by incorporating the perceptual quality of the decompressed-image, exposing an unavoidable trade-off between distortion and visual quality. Our work practices this very rationale by sacrificing image content (via pre-editing) for obtaining better looking compressed-decompressed images. The work by Agustsson et. al [2] offers a GAN-based learned compression algorithm that practically trades visual quality for distortion. While aiming for the same goal as our work, [2] replaces the whole compression-decompression process, whereas we insist on boosting available standard algorithms, such as JPEG, due to their massive availability and spread use.

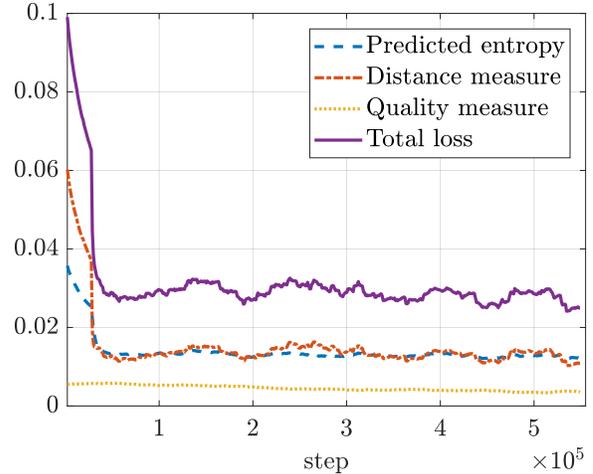


Figure 7: Our training loss components during gradient descent with JPEG quality factor in the range [8,25]. For better display, all losses are smoothed.

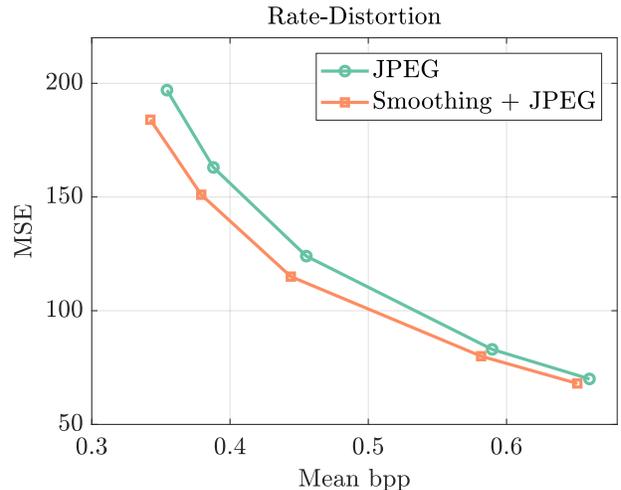


Figure 8: MSE vs. mean bit-rate for the Kodak dataset [19].

5. Experimental Results

In this section our results are discussed and compared to other methods. Our train and test are performed on a single Nvidia GPU V100 with 16GB RAM. At training, images are cropped to 480×640 , and testing is performed on the Kodak dataset [19]. We use the Adam optimizer [32] with learning rate set to 0.0001, and batch size as 1. The editing network is trained for 5×10^5 steps of stochastic gradient descent. Weights from the NIMA are kept fixed during training.

In order to train the STN and smoothing networks, we randomly sample the JPEG quality factor from a uniform distribution in the range [8, 25] at each step of the gradient descent. This allows our editing to be effective for a range of bit-rates. At test time, we compare our results with the baseline JPEG at comparable bit-rates. To compress a test

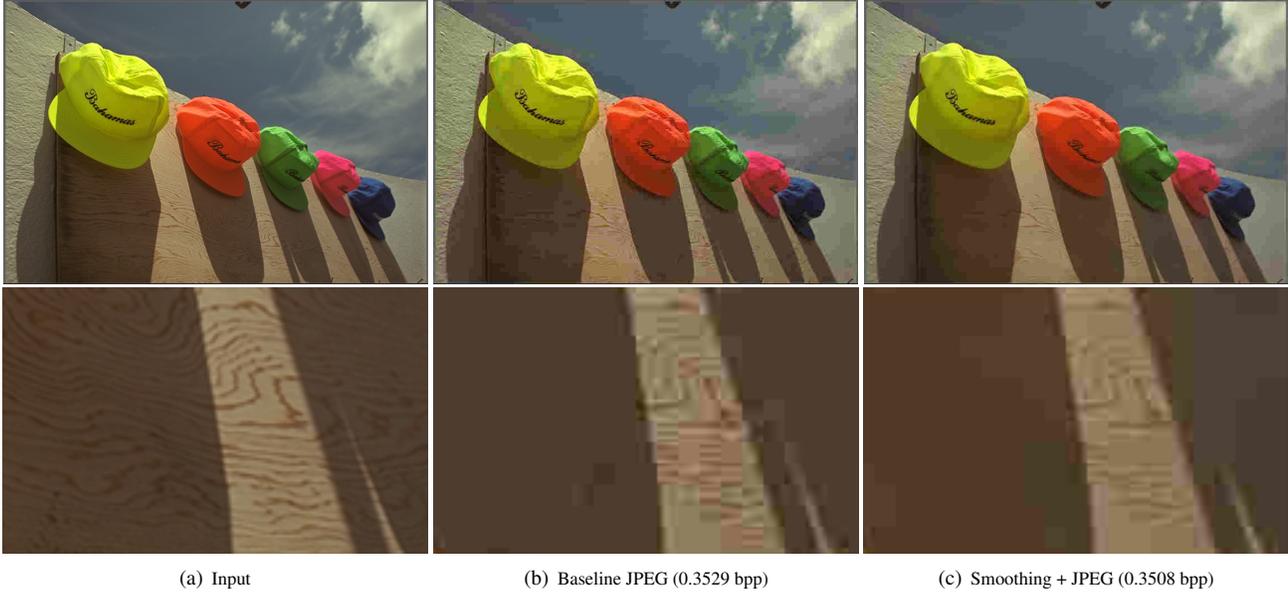


Figure 9: Compression performance with our smoothing network. Smoothing the image before compression leads to less blockiness and color artifacts.



Figure 10: Compression performance of the STN network. The STN applies local warps that lead to better detail preservation after compression.

image at various bit-rates, we adjust the JPEG quality factor to ensure that our result compresses with fewer bits.

Our weighted loss for cascaded smoother and STN are shown in Fig. 7. We select the NIMA weight to be $\lambda = 0.02$ and adjust the entropy prediction weight to $\mu = 1$. Our experiments suggest that the weighted predicted entropy and the distance measure should be close to each other. Also, as discussed in [46], the quality measure is most effective when its contribution is limited to a fraction of the total loss.

We trained the smoother and STN networks separately,

and then fine-tuned them jointly. The pre-trained smoothing is obtained by training with an L_2 distance measure $dist(\mathbf{x}_1, \mathbf{x}_2)$. Training images are augmented with random additive white Gaussian noise to enforce smoothing property in the resulting network. We randomly vary the standard deviation of the noise in the range $[0, 0.15]$ at each training step, and append the noise standard deviation and JPEG quality factor as extra channels to the input RGB image. Rate-distortion curves of the regular JPEG and the smoother content are shown in Fig. 8. As expected, the



Figure 11: Compression performance for applying smoothing and STN.

smoothing improves upon the baseline JPEG. Note that these results are obtained before fine-tuning the smoother with the STN. Examples of the smoother editing are shown in Fig. 9, where color degradation and blockiness artifacts are more visible in the baseline JPEG, compared to our results.

Results for training the STN network are shown in Fig. 10. Our editing of the input images allows to preserve structures and textures more effectively. The local deformations of the STN seem to make certain image textures more compressible. Note that this is a different behavior than the smoother’s effect .

We fine-tune both the smoother and STN networks jointly and present the results in Figs. 11 and 12. The cascade editor seems to present comparable details to the baseline, but with less JPEG artifacts.

We carried a human evaluation study to compare our proposed framework with baseline JPEG. We used Amazon Mechanical Turk with pairwise comparison for this task. We asked raters to select the image with better quality. We processed 24 images from Kodak dataset [19] with our smoothing and warping (STN) frameworks and compared them with their baseline JPEG counterparts at similar bit-rate. Comparisons are made by 20 human raters, and average percentage of the raters preference over baseline JPEG is reported in Fig. 13. As can be seen, both STN and our smoothing show perceptual preference of more than 50% for bit-rates smaller than 0.5 bpp. For higher bit-rates

our methods did not provide a statistically significant advantage over baseline. Also, we observed that smoothing consistently outperforms STN.

We conclude by referring to run time: We ran both our editors on an Intel Xeon CPU @ 3.5 GHz with 32 GB memory and 12 cores. We only measure timing of the pre-editing operation, as both methods use the JPEG encoder. The smoothing CNN and STN run in 1.7 sec and 1.2 sec on a 1 mega pixel image, respectively. Since our editors are based on convolutional neural networks, these running times can be further improved by GPU inference.

6. Conclusion

One of the main bottlenecks of low bit-rate JPEG compression is loss of textures and details and presence of visual artifacts. In this work we have proposed an end-to-end trainable manipulation framework that edits images before compression in order to mitigate these problems. Our CNN-based trained editors optimize for better perceptual quality, lower JPEG distortions and color degradation. The proposed image-editors are trained offline, avoiding the need for per-image optimization, or a post-processing on the decoder (client) side. Our future work will focus on extending this idea to other image compression standards, while seeking new ways to allow for more daring editing effects.



(a) Input (b) Baseline JPEG (0.4293 bpp) (c) Smoothing + STN + JPEG (0.4169 bpp)

Figure 12: Compression performance for applying smoothing and STN.

References

- [1] Requirements for still image coding using hevc. <http://mpeg.chiariglione.org/standards/mpeg-h/high-efficiency-video-coding/requirements-still-image-coding-using-hevc>, 2013.
- [2] Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool. Generative adversarial networks for extreme learned image compression. In *arXiv:1804.02958v3*, 2019.
- [3] Francois Alter, Sylvain Durand, and Jacques Froment. Adapted total variation for artifact free decompression of JPEG images. *Journal of Mathematical Imaging and Vision*, 23(2):199–211, 2005.
- [4] Amir Z Averbuch, Alon Schclar, and David L Donoho. Deblocking of block-transform compressed images using weighted sums of symmetrically aligned pixels. *IEEE Transactions on Image Processing*, 14(2):200–212, 2005.
- [5] J Ballé, V Laparra, and E P Simoncelli. End-to-end optimized image compression. In *Int'l. Conf. on Learning Representations (ICLR2017)*, Toulon, France, April 2017.
- [6] Sajjad Beygi, Shirin Jalali, Arian Maleki, and Urbashi Mitra. An efficient algorithm for compression-based compressed sensing. *Information and Inference: A Journal of the IMA*, 8(2):343–375, 2018.
- [7] Yochai Blau and Tomer Michaeli. Rethinking lossy compression: The rate-distortion-perception tradeoff. In *International Conf. on Machine Learning (ICML)*, 2019.
- [8] Sebastian Bosse, Dominique Maniry, Thomas Wiegand, and Wojciech Samek. A deep neural network for image quality assessment. In *IEEE International Conference on Image Processing (ICIP)*, pages 3773–3777, 2016.
- [9] Alfred M Bruckstein, Michael Elad, and Ron Kimmel. Down-scaling for better transform compression. *IEEE Transactions on Image Processing*, 12(9):1132–1144, 2003.
- [10] Lukas Cavigelli, Pascal Hager, and Luca Benini. Cas-cnn: A deep convolutional neural network for image compression artifact suppression. In *International Joint Conference on Neural Networks (IJCNN)*, 2017.

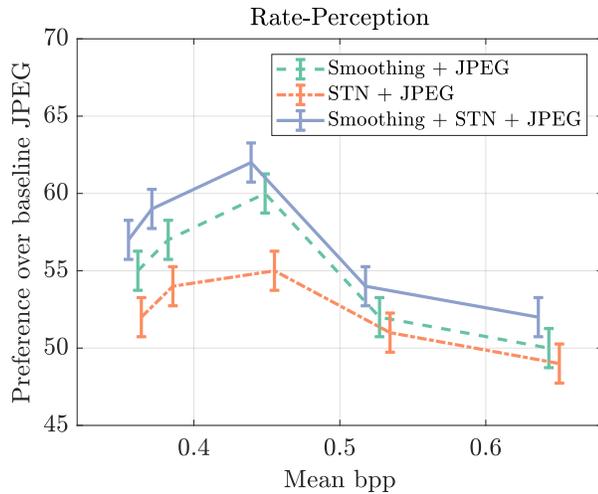


Figure 13: Percentage of human raters preference for pairwise comparison between our result and baseline JPEG. Each data point is an average of 480 ratings (24 Kodak images [19] and 20 human raters).

- [11] Tao Chen, Hong Ren Wu, and Bin Qiu. Adaptive postfiltering of transform coefficients for the reduction of blocking artifacts. *IEEE Transactions on Circuits and Systems for Video Technology*, 11(5):594–602, 2001.
- [12] Charilaos Christopoulos, Athanassios Skodras, and Touradj Ebrahimi. The JPEG2000 still image coding system: an overview. *IEEE Transactions on Consumer Electronics*, 46(4):1103–1127, 2000.
- [13] Yehuda Dar, Alfred M Bruckstein, Michael Elad, and Raja Giryes. Postprocessing of compressed images via sequential denoising. *IEEE Transactions on Image Processing*, 25(7):3044–3058, 2016.
- [14] Yehuda Dar, Michael Elad, and Alfred M Bruckstein. Optimized pre-compensating compression. *IEEE Transactions on Image Processing*, 27(10):4798–4809, 2018.
- [15] Yehuda Dar, Michael Elad, and Alfred M Bruckstein. System-aware compression. In *IEEE International Symposium on Information Theory (ISIT)*, pages 2226–2230, 2018.
- [16] Chao Dong, Deng Yubin, Change Loy Chen, and Xiaoou Tang. Compression artifacts reduction by a deep convolutional network. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.
- [17] Ke Du, Haiyun Han, and Gang Wang. A new algorithm for removing compression artifacts of wavelet-based image. In *IEEE International Conference on Computer Science and Automation Engineering*, volume 1, pages 336–340, 2011.
- [18] Ke Du, Jianming Lu, Hiroo Sekiya, and Takashi Yahagi. Post-processing for restoring edges and removing artifacts of low bit rates wavelet-based image. *IEEE Transactions on Electronics, Information and Systems*, 127(6):928–936, 2007.
- [19] Rich Franzen. Kodak lossless true color image suite. <http://r0k.us/graphics/kodak>, 1999.
- [20] Leonardo Galteri, Lorenzo Seidenari, marco Bertini, and Alberto Del Bimbo. Deep generative adversarial compression artifact removal. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4826–4835, 2017.
- [21] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016.
- [22] Maurizio Pintus Ginesu, Giaime and Daniele D. Giusto. Objective assessment of the webp image coding algorithm. *Signal Processing: Image Communication*, 27(8):867–874, 2012.
- [23] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315–323, 2011.
- [24] Samuel W. Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T. Barron, Florian Kainz, Jiawen Chen, and Marc Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 35(6), 2016.
- [25] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [26] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015.
- [27] Bin Jin, Maria V Ortiz Segovia, and Sabine Süsstrunk. Image aesthetic predictors based on weighted cnns. In *IEEE International Conference on Image Processing (ICIP)*, pages 2291–2295, 2016.
- [28] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, pages 694–711, 2016.
- [29] Rajan L Joshi, Majid Rabbani, and Margaret A Lepley. Comparison of multiple compression cycle performance for JPEG and JPEG 2000. In *Applications of Digital Image Processing XXIII*, volume 4115, pages 492–501, 2000.
- [30] Cheolkon Jung, Licheng Jiao, Hongtao Qi, and Tian Sun. Image deblocking via sparse representation. *Signal Processing: Image Communication*, 27(6):663–677, 2012.
- [31] Tomislav Kartalov, Zoran A Ivanovski, Ljupcho Panovski, and Lina J Karam. An adaptive pocs algorithm for compression artifacts removal. In *2007 9th International Symposium on Signal Processing and its Applications*, pages 1–4, 2007.
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [33] Ivan Kopilovic and Tams Szirnyi. Artifact reduction with diffusion preprocessing for image compression. *Optical Engineering*, 44(2):027003, 2005.
- [34] Younghee Kwon, Kwang In Kim, James Tompkin, Jin Hyung Kim, and Christian Theobalt. Efficient learning of image super-resolution and compression artifact removal with semi-local gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1792–1805, 2015.

- [35] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- [36] Weisi Lin and Li Dong. Adaptive downsampling to improve image compression at low bit rates. *IEEE Transactions on Image Processing*, 15(9):2513–2521, 2006.
- [37] Xin Lu, Zhe Lin, Xiaohui Shen, Radomir Mech, and James Z Wang. Deep multi-patch aggregation network for image style, aesthetics, and quality estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 990–998, 2015.
- [38] Long Mai, Hailin Jin, and Feng Liu. Composition-preserving deep photo aesthetics assessment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 497–506, 2016.
- [39] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing*, 21(12):4695–4708, 2012.
- [40] Munenori Oizumi. Preprocessing method for dct-based image-compression. *IEEE Transactions on Consumer Electronics*, 52(3):1021–1026, 2006.
- [41] Tamar Rott Shaham and Tomer Michaeli. Deformation aware image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2453–2462, 2018.
- [42] Mei-Yin Shen and C-C. Jay Kuo. Review of postprocessing techniques for compression artifact removal. *Journal of Visual Communication and Image Representation*, 9(1):2–14, 1998.
- [43] Richard Shin and Dawn Song. JPEG-resistant adversarial images. In *NIPS 2017 Workshop on Machine Learning and Computer Security*, 2017.
- [44] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [45] Jean-Luc Starck, Fionn Murtagh, Benoit Pirenne, and Miguel Albrecht. Astronomical image compression based on noise suppression. *Publications of the Astronomical Society of the Pacific*, 108(723), 1996.
- [46] Hossein Talebi and Peyman Milanfar. Learned perceptual image enhancement. In *2018 IEEE International Conference on Computational Photography (ICCP)*, pages 1–13. IEEE, 2018.
- [47] Hossein Talebi and Peyman Milanfar. NIMA: Neural image assessment. *IEEE Transactions on Image Processing*, 27(8):399–4011, 2018.
- [48] GA Triantafyllidis, D Sampson, D Tzovaras, and MG Strintzis. Blockiness reduction in JPEG coded images. In *2002 14th International Conference on Digital Signal Processing Proceedings. DSP 2002 (Cat. No. 02TH8628)*, volume 2, pages 1325–1328, 2002.
- [49] Yaakov Tsaig, Michael Elad, Peyman Milanfar, and Gene H Golub. Variable projection for near-optimal filtering in low bit-rate block coders. *IEEE Transactions on Circuits and Systems for Video Technology*, 15(1):154–160, 2005.
- [50] Florence Tushabe and M. H. F. Wilkinson. Image preprocessing for compression: Attribute filtering. In *Proceedings of International Conference on Signal Processing and Imaging Engineering (ICSPIE07)*, 2007.
- [51] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- [52] Gregory K Wallace. The JPEG still picture compression standard. *IEEE Transactions on Consumer Electronics*, 38(1):xviii–xxxiv, 1992.
- [53] Zhangyang Wang, Ding Liu, Shiyu Chang, Qing Ling, Yingzhen Yang, and Thomas S. Huang. D3: Deep dual-domain based fast restoration of JPEG-compressed images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [54] Pierre Weiss, Laure Blanc-Féraud, Thomas André, and Marc Antonini. Compression artifacts reduction using variational methods: Algorithms and experimental study. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1173–1176, 2008.
- [55] Xiaolin Wu, Xiangjun Zhang, and Xiaohan Wang. Low bit-rate image compression via adaptive down-sampling and constrained least squares upconversion. *IEEE Transactions on Image Processing*, 18(3):552–561, 2009.
- [56] Wufeng Xue, Lei Zhang, and Xuanqin Mou. Learning without human scores for blind image quality assessment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 995–1002, 2013.
- [57] Avideh Zakhor. Iterative procedures for reduction of blocking effects in transform image coding. *IEEE Transactions on Circuits and Systems for Video Technology*, 2(1):91–95, 1992.
- [58] Guangtao Zhai, Weisi Lin, Jianfei Cai, Xiaokang Yang, and Wenjun Zhang. Efficient quadtree based block-shift filtering for deblocking and deringing. *Journal of Visual Communication and Image Representation*, 20(8):595–607, 2009.
- [59] Xinfeng Zhang, Ruiqin Xiong, Xiaopeng Fan, Siwei Ma, and Wen Gao. Compression artifact reduction by overlapped-block transform coefficient estimation with block similarity. *IEEE Transactions on Image Processing*, 22(12):4613–4626, 2013.