Postprocessing of Compressed Images via Sequential Denoising

Yehuda Dar, Alfred M. Bruckstein, Michael Elad, Fellow, IEEE, and Raja Giryes

Abstract—In this paper, we propose a novel postprocessing technique for compression-artifact reduction. Our approach is based on posing this task as an inverse problem, with a regularization that leverages on existing state-of-the-art image denoising algorithms. We rely on the recently proposed Plugand-Play Prior framework, suggesting the solution of general inverse problems via alternating direction method of multipliers, leading to a sequence of Gaussian denoising steps. A key feature in our scheme is a linearization of the compression-decompression process, so as to get a formulation that can be optimized. In addition, we supply a thorough analysis of this linear approximation for several basic compression procedures. The proposed method is suitable for diverse compression techniques that rely on transform coding. In particular, we demonstrate impressive gains in image quality for several leading compression methods—JPEG, JPEG2000, and HEVC.

Index Terms—Lossy compression, postprocessing, deblocking, denoising, image restoration, Plug-and-Play Prior.

I. INTRODUCTION

B ANDWIDTH and memory constraints play a crucial role in transmission and storage systems. Various compression methods are available in order to meet severe constraints on the bit-cost in data representation. While some applications require perfect reconstruction, some may tolerate inaccuracies and can benefit from a reduced representation-cost. The latter approach is known as lossy compression and is widely used for representing a signal under bit-budget constraints while allowing some errors in recovery. Accordingly, a variety of techniques were standardized over the years for the lossy compression of acoustic and visual signals.

Since lossy compression allows discrepancies between the original and the reconstructed signals, the differences being intentionally used in tradeoffs between bit-rate and quality. The nature of the created artifacts depends on the compression architecture. For example, block-based image compression techniques suffer from blockiness effects that increase and degrade the reconstruction as the bit-rate is reduced.

As artifacts are inherent in the lossy compression of signals, a great number of artifact-reduction techniques were proposed

Manuscript received October 27, 2015; revised March 17, 2016; accepted April 12, 2016. Date of publication April 27, 2016; date of current version May 17, 2016. This work was supported by the European Research Council through European Union Seventh Framework Program under Grant 320649. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Chang-Su Kim.

Y. Dar, A. M. Bruckstein, and M. Elad are with the Department of Computer Science, Technion–Israel Institute of Technology, Haifa 3200003, Israel (e-mail: ydar@cs.technion.ac.il; freddy@cs.technion.ac.il; elad@cs.technion.ac.il).

R. Giryes is with the School of Electrical Engineering, Tel Aviv University, Tel Aviv 6997801, Israel (e-mail: raja@tauex.tau.ac.il).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIP.2016.2558825

over the years (e.g., [1]-[20] for image compression). These methods usually focus on specific signal types (e.g., image, video or audio) and sometimes even on specific artifacts corresponding to certain compression designs (e.g., deblocking procedures for images). Common image compression techniques rely on transform-coding, where image blocks are transformed, and the resultant transformcoefficients are quantized according to their relative importance. The prominent artifacts of this architecture are [21]: blockiness due to the separate treatment of non-overlapping blocks; ringing caused by the effective elimination of high frequency components, expressed as contours spreading along sharp edges; and blurring that results from high-frequency information loss. Postprocessing of compressed images are subcategorized into two approaches [21]: enhancement of the deteriorated signal by smoothing its artifacts (e.g., [1], [3]), and restoration of the original signal samples (e.g., [2], [14]).

In this work we propose a novel postprocessing technique for compression artifact reduction by a regularized restoration of the original (precompressed) signal. Specifically, we formulate the compression postprocessing procedure as a regularized inverse-problem for estimating the original signal given its reconstructed form. We also approximate the (nonlinear!) compression-decompression process by a linear operator, so as to obtain a tractable inverse problem formulation. The intriguing approach of locally linearizing the non-differentiable compression procedures is carefully analyzed, in order to utilize it properly. Whereas many studies focus on corrections of specific artifacts (e.g., image deblocking techniques [1], [3], [5], [12]), our approach attempts to generally restore the signal and thus implicitly repairs multiple artifacts. The major strength of our method comes from the regularization used, as we next explain.

Afonso *et al.* [22] proposed to efficiently solve regularized inverse-problems in image processing using the alternating direction method of multipliers (ADMM) [23]. Their approach decouples the inversion and the regularization parts of the optimization problem, which is in turn iteratively solved. Venkatakrishnan *et al.* [24] further developed the use of the ADMM by showing an equivalence between the regularization step and denoising optimization problems. Their framework, called "Plug-and-Play Priors", is flexible, proposing the replacement of the regularization step by a general-purpose Gaussian image denoiser.

In this work we propose a compression postprocessing algorithm by employing the Plug-and-Play Priors framework. Furthermore, as denoising algorithms relying on sparse models were found to be highly effective ones (e.g., K-SVD [25], [26], BM3D [27]), we utilize a leading denoiser from this category.

1057-7149 © 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

The Plug-and-Play Priors framework was proposed for general inverse problems and was specifically demonstrated for reconstruction of tomographic images. The novelty of our work with respect to the original Plug-and-Play approach is that we apply it for the task of compression-artifact reduction. Moreover, we utilize it to address an inverse-problem for a forward-model that is non-linear and non-differentiable.

Since we propose a method of postprocessing for a variety of lossy-compression techniques, the algorithm and its analysis are dealt within an abstract and general setting. Following that, a thorough demonstration for image compression is provided. Specifically, we show results for the leading image compression standards: JPEG [28], JPEG2000 [29] and the still-image profile of the HEVC [30], [31], offering the state-of-the-art performance [32]. While these three compression methods rely on a block-based architecture and a transform-coding approach, they differ as follows: JPEG operates on 8×8 blocks and applies a discrete cosine transform (DCT); JPEG2000 works on large blocks (tiles) of at least 128×128 pixels and utilizes a discrete wavelet transform (DWT); in HEVC-stills the image is split into coding blocks that are further partitioned using a quadtree structure, then intra-prediction is performed and transform coding is applied on the prediction residuals (where the transform is mainly integer-approximations of the DCT at various sizes). Our method is evaluated for a diversified set of compression algorithms that span the range of the contemporary coding concepts. Moreover, our postprocessing technique achieves significant gains and usually outperforms the cutting-edge methods for the examined compression standards.

This paper is organized as follows. In section II the proposed postprocessing method is presented. In section III, the compression linearization is mathematically analyzed for simplified cases of quantization and transform coding. Section IV presents image-compression experimental results and compares them to those of competitive techniques. Section V concludes this paper.

II. THE PROPOSED POSTPROCESSING STRATEGY

A. Problem Formulation Using ADMM

Let us consider a signal $\mathbf{x} \in \mathbb{R}^N$ that undergoes a compression-decompression procedure, $C : \mathbb{R}^N \to \mathbb{R}^N$, resulting in the reconstructed signal $\mathbf{y} = C(\mathbf{x})$. For lossy compression methods an error is introduced at a size that depends on the bit-budget, the specific-signal characteristics, and the compression algorithm. We aim at restoring the precompressed signal \mathbf{x} from the reconstructed \mathbf{y} using the following regularized inverse-problem:

$$\hat{\mathbf{x}} = \operatorname*{arg\,min}_{\mathbf{y}} \|\mathbf{y} - C(\mathbf{x})\|_2^2 + \beta s(\mathbf{x}), \qquad (1)$$

where $s(\cdot)$ is a regularizer, which can be associated with a given Gaussian denoiser, weighted by the parameter β . For example, assuming that the image is piecewise constant promotes the utilization of the popular total-variation regularizer, $s(\mathbf{x}) = ||\mathbf{x}||_{TV}$ [33].

One should note that \mathbf{y} and $C(\mathbf{x})$ are two signals reconstructed from compression, and therefore, the fidelity term

in Equation (1) expresses their distance. Notice that this is substantially different from $\|\mathbf{y} - \mathbf{x}\|_2^2$ – whereas the latter has compression artifacts as error, the one we deploy represents a milder distortion. Throughout this paper we shall assume for simplicity that the distortion between the two reconstructions, **y** and *C* (**x**), is modeled as a white additive Gaussian noise, leading to the ℓ_2 term used here. We should note, however, that our scheme could be improved by using a better modeling of the reconstructed-signal error, such as an ℓ_{∞} on the transform coefficients w.r.t. the quantization step-size (in the case of transform coding).

Similar to [22] and [24], we develop an iterative algorithm for the solution of (1). We start by applying variable splitting that yields the following equivalent form of (1):

$$\min_{\mathbf{x}, \mathbf{v}} \|\mathbf{y} - C(\mathbf{x})\|_2^2 + \beta s(\mathbf{v})$$

subject to $\mathbf{x} = \mathbf{v}$, (2)

where $\mathbf{v} \in \mathbb{R}^N$ is an additional vector due to the split. The constrained problem (2) is addressed by forming an augmented Lagrangian and its corresponding iterative solution (of its scaled version) via the method of multipliers [23, Ch. 2], where the *i*th iteration consists of

$$\begin{aligned} \left(\hat{\mathbf{x}}_{i}, \hat{\mathbf{v}}_{i} \right) &= \operatorname*{arg\,min}_{\mathbf{x}, \mathbf{v}} \|\mathbf{y} - C\left(\mathbf{x}\right)\|_{2}^{2} + \beta s\left(\mathbf{v}\right) \\ &+ \frac{\lambda}{2} \|\mathbf{x} - \mathbf{v} + \mathbf{u}_{i}\|_{2}^{2} \\ \mathbf{u}_{i+1} &= \mathbf{u}_{i} + \left(\hat{\mathbf{x}}_{i} - \hat{\mathbf{v}}_{i} \right). \end{aligned}$$
(3)

Here $\mathbf{u}_i \in \mathbb{R}^n$ is the scaled dual-variable and λ is an auxiliary parameter, both introduced in the Lagrangian.

Please note the following notation remark for a general vector **u**. First, \mathbf{u}_i stands for vector **u** in the i^{th} iteration. On the other hand, u_j represents the j^{th} component (a scalar) of the vector **u**. Finally, $\mathbf{u}_i^{(j)}$ denotes the j^{th} element of the vector \mathbf{u}_i .

Approximating the joint optimization of **x** and **v** in (3), using one iteration of alternating minimization, results in the iterative solution in the ADMM form, where the i^{th} iteration consists of

$$\hat{\mathbf{x}}_{i} = \operatorname*{arg\,min}_{\mathbf{x}} \|\mathbf{y} - C\left(\mathbf{x}\right)\|_{2}^{2} + \frac{\lambda}{2} \|\mathbf{x} - \tilde{\mathbf{x}}_{i}\|_{2}^{2}$$
(4)

$$\hat{\mathbf{v}}_{i} = \arg\min_{\mathbf{v}} \frac{\lambda}{2} \|\mathbf{v} - \tilde{\mathbf{v}}_{i}\|_{2}^{2} + \beta s(\mathbf{v})$$
(5)

$$\mathbf{u}_{i+1} = \mathbf{u}_i + \left(\hat{\mathbf{x}}_i - \hat{\mathbf{v}}_i\right). \tag{6}$$

Here $\tilde{\mathbf{x}}_i = \hat{\mathbf{v}}_{i-1} - \mathbf{u}_i$ and $\tilde{\mathbf{v}}_i = \hat{\mathbf{x}}_i + \mathbf{u}_i$.

The regularization step (5) is of the form of a Gaussian denoising optimization-problem (of a noise level determined by β/λ) and therefore can be viewed as applying a denoising algorithm to the signal $\tilde{\mathbf{v}}_i$. More specifically, this corresponds to assuming that $\tilde{\mathbf{v}}_i = \mathbf{v} + \mathbf{w}$, where \mathbf{w} is an i.i.d zero-mean Gaussian vector with variance $1/\lambda$ (and a corresponding distribution function denoted as $p_w(\mathbf{w})$). In addition, \mathbf{v} is assumed to be drawn from a distribution $p_s(\mathbf{v})$ that is proportional to $\exp(-\beta s(\mathbf{x}))$. Then, the Maximum A-Posteriori (MAP) estimator of \mathbf{v} from its (white Gaussian) noisy version $\tilde{\mathbf{v}}_i$ is formed as

$$\hat{\mathbf{v}}_i = \arg\max_{\mathbf{v}} \log p_w(\tilde{\mathbf{v}}_i - \mathbf{v}) + \log p_s(\mathbf{v}), \tag{7}$$

which for the above defined distribution functions, $p_w(\cdot)$ and $p_s(\cdot)$, is equivalent to (5) and, thus, establishes the latter as a Gaussian denoising procedure. Indeed, the Plug-and-Play Priors framework [24] suggests exactly this strategy, replacing (5) with an independent denoiser; even one that does not explicitly have in its formulation a minimization problem of the form of (5). The deployment of a favorable denoiser introduces valuable practical benefits to the design of the proposed postprocessing procedure, and yields a powerful generic method.

B. Linear Approximation of the Compression-Decompression Procedure

Due to the high nonlinearity of $C(\mathbf{x})$, we further simplify the forward-model step (4) using a first-order Taylor approximation of the compression-decompression function around $\hat{\mathbf{x}}_{i-1}$, i.e.,

$$C_{lin}\left(\mathbf{x}\right) = C\left(\hat{\mathbf{x}}_{i-1}\right) + \left.\frac{dC\left(\mathbf{z}\right)}{d\mathbf{z}}\right|_{\mathbf{z}=\hat{\mathbf{x}}_{i-1}} \cdot \left(\mathbf{x}-\hat{\mathbf{x}}_{i-1}\right) \qquad (8)$$

where $\frac{dC(\mathbf{z})}{d\mathbf{z}}\Big|_{\mathbf{z}=\hat{\mathbf{x}}_{i-1}}$ is the $N \times N$ Jacobian matrix of the compression-decompression at the point $\hat{\mathbf{x}}_{i-1}$.

Since the approximation of the Jacobian, $\frac{dC(z)}{dz}$, deeply influences the restoration result and the computational cost, this is a quite delicate task. First, *C* is a non-linear and even non-differentiable function as the compression often relies on quantization and/or thresholding. Second, we provide here a generic technique, and therefore do not explicitly consider the compression-decompression formulation.

Theoretically, the fact that C is non-differentiable would prevent us from using its Jacobian for the optimization. However, as its Jacobian has only a finite number of singularities, in practice we can rely on it as is done in other fields, e.g., in the training of neural networks that are composed of concatenations of non-differentiable non-linear operations [34].

For calculating the entries of the Jacobian, we rely on the standard definition of the derivative, assuming that *C* is locally linear. We justify this approach in the next section. As we might be approximating the derivative in the neighborhood of a non-differential point, we take several step-sizes in the calculation of the derivative and average over all of them. This leads to the following approximation to the k^{th} column of the Jacobian:

$$\frac{dC(\mathbf{z})}{dz_k} = \frac{1}{|S_{\delta}|} \sum_{\delta \in S_{\delta}} \frac{C(\mathbf{z} + \delta \cdot \mathbf{e}_k) - C(\mathbf{z} - \delta \cdot \mathbf{e}_k)}{2\delta}, \quad (9)$$

where \mathbf{e}_k is the k^{th} standard direction vector, and S_{δ} is a set of step lengths for approximating the derivative using the standard definition (the set size is denoted as $|S_{\delta}|$).

Due to the high nonlinearity of *C*, the linear approximation (8) is reasonable in a small neighborhood around the approximating point $\hat{\mathbf{x}}_{i-1}$. Accordingly, we further constrain the distance of the solution from the linear-approximation

Algorithm 1 The Proposed Postprocessing Method

1: $\hat{\mathbf{x}}_0 = \mathbf{y}$, $\hat{\mathbf{v}}_0 = \mathbf{y}$ 2: i = 1, $\mathbf{u}_1 = \mathbf{0}$ 3: repeat 4: Approximate $C_{lin}(\cdot)$ around $\hat{\mathbf{x}}_{i-1}$ using (8) and (9) 5: $\tilde{\mathbf{x}}_i = \hat{\mathbf{v}}_{i-1} - \mathbf{u}_i$ $\hat{\mathbf{x}}_{i} = \underset{\mathbf{x}}{\arg\min} \|\mathbf{y} - C_{lin}(\mathbf{x})\|_{2}^{2} + \frac{\lambda}{2} \|\mathbf{x} - \tilde{\mathbf{x}}_{i}\|_{2}^{2}$ $\mathbf{\tilde{v}}_{i} = \mathbf{\hat{x}}_{i} + \mathbf{u}_{i}$ $\mathbf{\tilde{v}}_{i} = \mathbf{\hat{x}}_{i} + \mathbf{u}_{i}$ 6: 7: $\hat{\mathbf{v}}_{i} = \text{Denoise}_{\beta/\lambda} \left(\tilde{\mathbf{v}}_{i} \right)$ 8: $\mathbf{u}_{i+1} = \mathbf{u}_i + (\hat{\mathbf{x}}_i - \hat{\mathbf{v}}_i)$ 9: $i \leftarrow i + 1$ 10: 11: until stopping criterion is satisfied

point by modifying (4) to

$$\hat{\mathbf{x}}_{i} = \underset{\mathbf{x}}{\operatorname{arg\,min}} \|\mathbf{y} - C_{lin}(\mathbf{x})\|_{2}^{2} + \frac{\lambda}{2} \|\mathbf{x} - \tilde{\mathbf{x}}_{i}\|_{2}^{2} + \mu \|\mathbf{x} - \hat{\mathbf{x}}_{i-1}\|_{2}^{2}.$$
(10)

The proposed generic method is summarized in Algorithm 1.

III. LINEAR APPROXIMATION-A CLOSER LOOK

The wide variety of lossy-compression methods yield a range of diverse compression performances and features. These often rely on the fundamental procedure of quantization, which enables to trade-off representation-precision and cost (in bits). The quantization concept is employed in various forms, e.g., as a scalar/vector operation, using uniform/ non-uniform representation levels, and also in the extreme case of thresholding where some data elements are completely discarded (and the remaining are regularly quantized). Furthermore, the statistical properties of the data affect the quantizer performance, and indeed, the prevalent transformcoding concept first considers the data in different orthogonal basis that enables more efficient quantization.

In this section we study the linear approximation of the scalar-quantization procedure, starting at its use for a single variable and proceeding to transform coding of vectors, where the transform-domain coefficients are independently quantized. The analysis provided here sheds some clarifying light on our linearization strategy that is generically applied in the proposed technique to more complicated compression methods.

A. Local Linear-Approximation of a Quantizer

Let us consider a general scalar quantization function q(x) that maps the real-valued input x onto a discrete set of realvalued representation levels. As q(x) is a non-differentiable function we examine its linear approximation around the point x_0 in a limited interval defined by δ as

$$\eta(x_0, \delta) = [x_0 - \delta, x_0 + \delta]^{-1}$$
(11)

¹Here we mathematically study the problem for a given quantization function that is used for calculating the approximation, and therefore a single δ value is sufficient. However, in our generic algorithm we empirically utilize a set of δ values in Equation (9) since the compression function is unknown.

The studied approximation takes the general linear form of

$$\tilde{q}(x) = ax + b \tag{12}$$

where *a* and *b* are the linearization parameters. The approximation in (12) introduces an error that can be measured via the Mean-Squared-Error (MSE) over the local interval $\eta(x_0, \delta)^2$:

$$LMSE^{SQ}(a,b;\eta(x_0,\delta)) \triangleq \frac{1}{2\delta} \int_{x_0-\delta}^{x_0+\delta} (q(x) - \tilde{q}(x))^2 dx$$
(13)

Substituting (12) in (13), then demanding parameter optimality by

$$\frac{\partial}{\partial a}LMSE^{SQ}(a,b;\eta(x_0,\delta)) = 0$$
$$\frac{\partial}{\partial b}LMSE^{SQ}(a,b;\eta(x_0,\delta)) = 0, \qquad (14)$$

leads to the following optimal parameters:

$$a^{*} = \frac{3}{\delta^{2}} \left(L_{a} - L_{b} x_{0} \right)$$
(15)

$$b^* = L_b - \frac{3x_0}{\delta^2} \left(L_a - L_b x_0 \right)$$
(16)

where we defined

$$L_{a} \triangleq \frac{1}{2\delta} \int_{x_{0}-\delta}^{x_{0}+\delta} xq(x) dx$$
(17)

$$L_b \triangleq \frac{1}{2\delta} \int_{x_0-\delta}^{x_0+\delta} q(x) \, dx. \tag{18}$$

To better understand the values of these parameters, we shall consider several simple cases.

B. The Case of Two-Level Quantization

We start by studying the elementary two-level quantizer that takes the form of a step function (Fig. 1) as follows:

$$q_2(x) = \begin{cases} -1/2, & \text{for } x \le 0\\ 1/2, & \text{for } x > 0 \end{cases}$$
(19)

where the two output levels, $r_0 = -1/2$ and $r_1 = 1/2$, are assigned according to the input sign. This canonic form is useful to our discussion here, since it is an asymmetric function around the origin, and thus, will simplify the mathematical analysis. Nevertheless, the form in (19) can be extended to any two-level quantizer using shifts and scaling that adjust the steplocation and the two representation levels. Accordingly, the results in this section are easily extended, e.g., by considering the quadratic effect of the step-size scaling on the local MSE.



Fig. 1. Examples of scalar quantizers. (a) Two-level. (b) Uniform.

When the local interval is completely contained within a single decision region, i.e. $\eta(x_0, \delta) \subset [-\infty, 0]$ or $\eta(x_0, \delta) \subset (0, \infty]$, then $q_2(x)$ is locally fixed on r_0 or r_1 , respectively, and therefore

$$L_a = \frac{1}{2\delta} \int_{x_0 - \delta}^{x_0 + \delta} x r_i dx = x_0 r_i$$
(20)

$$L_b = \frac{1}{2\delta} \int_{x_0 - \delta}^{x_0 + \delta} r_i dx = r_i$$
(21)

for the respective $i \in \{0, 1\}$. Then setting (20) and (21) in (15) and (16), respectively, induces the optimal values $a^* = 0$ and $b^* = r_i$, that of course accurately represent the locally flat function with a corresponding zero local-MSE.

Now we turn to the more interesting case where the local interval spans over the two decision regions, i.e., $x_0 - \delta < 0$ and $x_0 + \delta > 0$. Calculating again the optimal parameter set (15)-(16) for this scenario requires to decompose the integrals (17)-(18) to the two decision regions, yielding the following optimal linearization parameters:

$$a^* = \frac{3}{4\delta} \left(1 - \left(\frac{x_0}{\delta}\right)^2 \right) \tag{22}$$

$$b^* = \frac{3x_0}{4\delta} \left(\left(\frac{x_0}{\delta}\right)^2 - \frac{1}{3} \right) \tag{23}$$

and, using (13), the corresponding error (for $\delta > |x_0|$) is³

L

$$MSE^{SQ}(a^*, b^*; \eta(x_0, \delta)) = \frac{1}{16} \left(1 + 3\left(\frac{x_0}{\delta}\right)^2\right) \left(1 - \left(\frac{x_0}{\delta}\right)^2\right).$$
(24)

One should note that on the limit of the global linear approximation, i.e., when $\delta \rightarrow \infty$, the optimal parameters are

$$\lim_{\delta \to \infty} a^* = 0 \tag{25}$$

$$\lim_{\delta \to \infty} b^* = 0. \tag{26}$$

This asymptotic fitting to a constant-valued function is also expressed in the numerical results in Fig. 2a-2b.

Let us study the optimal approximation for the non-trivial case of $\delta > |x_0|$. First, we notice that the error tends to zero as δ gets closer to $|x_0|$. Second, The maximal error is obtained for $\delta = \sqrt{3}|x_0|$ and its value is $\frac{1}{12}$. Moreover, for approximation around the non-differentiable point, i.e. $x_0 = 0$, the error is

³Recall that for $\delta < |x_0|$ the approximation local-MSE is zero.

²We study the linearization error as a function of the approximation-interval size, which is determined by δ . Clearly, by setting a sufficiently small δ we get a zero approximation-error as we shall see hereafter. However, note that the linearization error is not the only factor to consider for the selection of δ . Therefore, we should bear in mind throughout the following derivation that we do not present here an explicit method for selecting the value of δ but an analysis of the local approximation-error of the quantizer as a function of δ . Nevertheless, this mathematical analysis demonstrates the important principles of linear approximation of quantizers and motivates the algorithmic design and experimental settings that are presented in the following sections.



Fig. 2. Optimal linearization of a normalized two-level quantizer as function of the local interval center (x_0) and length (δ). (a) a^* . (b) b^* . (c) Local MSE.

a constant and, therefore, independent of δ . This interesting observation is a special case of a more general behavior where a constant error value is achieved for any (x_0, δ) pair that is on the line $\delta = c|x_0|$ for some $c \in [0, \infty)$. This constant local-MSE is due to the fixed ratio between the lengths of the subintervals $[x_0 - \delta, 0]$ and $[0, x_0 + \delta]$, determining the optimal approximation in this case. The latter analysis is clearly exhibited in the numerical results in Fig. 2.

The numerical results also demonstrate the following behavior of the approximation as function of δ . At the beginning, the solution gradually considers the step by having an increasingly steeper slope, then, the approximation begins to approach the asymptotic solution of a flat line. It is also observed that the approximation is useful (in terms of relatively low error) when the interval size tends to be the minimal that contains the discontinuity point, located here at 0. Furthermore, in some sense, finding the best interval for approximating around $x_0 \neq 0$ is like measuring the distance of x_0 from the step.

C. The Case of Multi-Level Uniform Quantization

Let us extend the above analysis to a multi-level uniform quantizer in the mid-riser form [35, p. 137] (Fig. 1b):

$$q_u(x) = \lfloor x \rfloor + \frac{1}{2}.$$
 (27)

Here the quantization step is of unit length, and accordingly the i^{th} decision region, $[d_i, d_{i+1}) = [i, i+1)$, maps the input to the i^{th} representation level $r_i = i + \frac{1}{2}$. Note that *i* is an integer that may be positive or negative. As in the previous case, this normalized quantizer form yields a simplified analysis that is, however, extendable to any uniform quantizer by shifts and scaling.

The optimal local linear approximation for this uniform quantizer is obtained by calculating (15)-(16) for the formula in (27). Again, the solution depends on the interval layout. In the simplest case, the considered interval is completely contained within a single decision region, i.e., $\eta(x_0, \delta) \subset [d_i, d_{i+1}]$ for some *i*. Here $q(x) = r_i$ for any $x \in \eta(x_0, \delta)$. Clearly, the corresponding discussion for the two-level quantizer (see section III-B) also holds here, meaning that $a^* = 0$ and $b^* = r_i$ with a zero local-MSE.

Another scenario that coincides with the two-level quantizer is when $x_0 - \delta \in [d_{i-1}, d_i]$ and $x_0 + \delta \in [d_i, d_{i+1}]$, i.e., the interval is spread over only two adjacent decision regions. Indeed, the optimal parameters here are obtained by appropriately shifting the results in (22)-(23). However, note that the multi-level quantizer has two levels only for $\delta < \min \{d_{i+1} - x_0, x_0 - d_{i-1}\} < 1$.

Now we proceed to the main case, where the approximation interval spans over more than two decision regions, i.e., $x_0 - \delta \in [d_i, d_{i+1}]$ and $x_0 + \delta \in [d_j, d_{j+1}]$ for j - i > 1. First, we express the uniform quantization function as a sum of shifted two-level quantizers:

$$q_u(x) = \sum_{\tau = -\infty}^{\infty} q_2(x - \tau),$$
 (28)

where $q_2(\cdot)$ was defined in (19). Then, using (28) we can develop (17)-(18) to the following forms:

$$L_a^u = \sum_{\tau = -\infty}^{\infty} L_a^{\tau} \tag{29}$$

$$L_b^u = \sum_{\tau = -\infty}^{\infty} L_b^{\tau},\tag{30}$$

where L_a^{τ} and L_b^{τ} are the corresponding values for the two-level quantizer $q_2(x - \tau)$. These allow us to write the optimal linearization parameters of the uniform quantizer as the summation of the optimal parameters of the shifted two-level quantizers, namely

$$a_u^* = \sum_{\tau = -\infty}^{\infty} a_\tau^* \tag{31}$$

$$b_u^* = \sum_{\tau = -\infty}^{\infty} b_\tau^*, \tag{32}$$

where a_{τ}^{*} and b_{τ}^{*} are the optimal linearization parameters for $q_2(x - \tau)$ and are obtainable by shifting the expressions in (22)-(23). This analytic relation between the linearization of the uniform and the two-level quantizers is clearly exhibited in the numerical results (see Fig. 3) in the form of a periodic structure.

The numerical calculations (Fig. 3) also show convergence to the global approximation parameters

$$\lim_{\delta \to \infty} a_u^* = 1 \tag{33}$$

$$\lim_{\delta \to \infty} b_u^* = 0, \tag{34}$$



Fig. 3. Optimal linearization of a normalized uniform quantizer as function of the local interval center (x_0) and length (δ) . (a) a_u^* . (b) b_u^* . (c) Local MSE.

which imply $\lim_{x \to \infty} C_{lin}(x) = x$. In order to explain the results in Fig. 3, we return to the interpretation of a multi-level quantizer as a sum of shifted two-level quantizers (as expressed in Eq. (28)). First, examining the case of approximation around a decision level, shows that at each point of $\delta = k$ (for integer k values), two additional representation levels are included in the approximation (one on each side of the interval) and affect the optimal approximation. Comparing Fig. 3 to Fig. 2 reveals that the effect of each of these added representation levels is like approximating a two-level quantizer around a point that differs from its threshold level. Evaluating the approximation around a point that is not a decision level (see Fig. 3 while considering non-integer x_0 values) extends the previous behavior by combining two unsynchronized periodic patterns, each of them stems from a recurrent addition of representation levels from a different side.

The MSE plot (Fig. 3c) shows that, for nontrivial intervals that contain at least one non-differentiable point, the minimal MSE is obtained for approximation over a small interval that includes only the nearest decision level. This somewhat resembles the underlying principle of the dithering procedure [36], where the points within a quantization-cell are differentiated by an added noise that statistically maps them to neighboring cells according to their relative proximity. Moreover, maximal MSE of 0.106 is obtained for $\delta = 0.67$ and $x_0 = \frac{1}{2} + i$ (for $i = 0, \pm 1, \pm 2, ...$), where only the two adjacent non-differentiable points affect the linearization. This can also be shown analytically by setting the decompositions in (28) and (31)-(32) into (13), resulting in

$$LMSE_{u}^{SQ}(a^{*}, b^{*}; \eta(x_{0}, \delta)) = \sum_{\tau=-\infty}^{\infty} LMSE_{\tau}^{SQ}(a_{\tau}^{*}, b_{\tau}^{*}; \eta(x_{0}, \delta)) + \frac{1}{2\delta} \sum_{\tau, \nu=-\infty}^{\infty} \int_{x_{0}-\delta}^{x_{0}+\delta} (q_{2}(x-\tau) - a_{\tau}^{*}x - b_{\tau}^{*}) \times (q_{2}(x-\nu) - a_{\nu}^{*}x - b_{\nu}^{*}) dx, \quad (35)$$

where $LMSE_{\tau}^{SQ}(a_{\tau}^*, b_{\tau}^*; \eta(x_0, \delta))$ is the optimal LMSE for $q_2(x - \tau)$ as available by shifting the expression in (24).

D. Transform Coding

We now turn to generalize the discussion to compression of multidimensional signals by considering the widely used concept of transform coding, where scalar quantization is applied in the transform domain. We examine coding of an N-length signal vector using a unitary transform, that can be formulated as the vector-valued function

$$C\left(\mathbf{x}\right) = \mathbf{U}Q\left(\mathbf{U}^{T}\mathbf{x}\right),\tag{36}$$

where **x** is the $N \times 1$ signal to compress, **U** is an $N \times N$ unitary matrix, and $Q(\cdot)$ is a vector-valued quantization function that scalarly quantizes the input components, i.e.,

$$Q\left(\mathbf{x}\right) = \begin{bmatrix} q\left(x_{1}\right) \\ \vdots \\ q\left(x_{N}\right) \end{bmatrix}$$
(37)

where $q(\cdot)$ is a single-variable scalar quantization function as studied above, and x_i is the *i*th component of the vector **x**. Moreover, as the last definition exhibits, the discussion is simplified by assuming identical quantization rules to all vector components.

As scalar quantization is a building block of the transform coding procedure (36), it imposes its non-differentiable nature on $C(\mathbf{x})$. Let us consider the linear approximation of $C(\mathbf{x})$ around the point $\mathbf{x}_0 \in \mathbb{R}^N$ in a limited neighborhood of a high-dimensional cube defined by δ as

$$\eta(\mathbf{x}_0, \delta) = \{ \mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_0\|_{\infty} \le \delta \}.$$
(38)

The approximation takes the general multidimensional linear form of

$$\tilde{C}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b} \tag{39}$$

where $\mathbf{A} \in \mathbb{R}^{N \times N}$ and $\mathbf{b} \in \mathbb{R}^N$ are the linearization parameters. The local MSE of approximating the transform-coding procedure around \mathbf{x}_0 is defined as

$$LMSE^{TC}(\mathbf{A}, \mathbf{b}; \eta(\mathbf{x}_{0}, \delta)) \triangleq \frac{1}{|\eta(\mathbf{x}_{0}, \delta)|} \int_{\eta(\mathbf{x}_{0}, \delta)} \|C(\mathbf{x}) - \tilde{C}(\mathbf{x})\|_{2}^{2} d\mathbf{x}$$
(40)

By substituting (39) in (40) and using the energy-preservation property of unitary transforms, we get the equivalent error expression in the transform-domain

$$\frac{1}{|\hat{\eta}(\mathbf{x}_0,\delta)|} \int_{\hat{\eta}(\mathbf{x}_0,\delta)} \left\| Q(\hat{\mathbf{x}}) - \hat{\mathbf{A}}\hat{\mathbf{x}} - \hat{\mathbf{b}} \right\|_2^2 d\hat{\mathbf{x}}$$
(41)

where

$$\hat{\mathbf{x}} = \mathbf{U}^{T} \mathbf{x}$$

$$\hat{\mathbf{x}}_{0} = \mathbf{U}^{T} \mathbf{x}_{0}$$

$$\hat{\mathbf{A}} = \mathbf{U}^{T} \mathbf{A} \mathbf{U}$$

$$\hat{\mathbf{b}} = \mathbf{U}^{T} \mathbf{b}$$

$$|\hat{\eta}(\mathbf{x}_{0}, \delta)| = |\eta(\mathbf{x}_{0}, \delta)|. \qquad (42)$$

and the rotated approximation area (or volume), $\hat{\eta}$, is defined around $\hat{\mathbf{x}}_0$ and may not have sides that are aligned with the axes.

Let us generally define the local linearization error for compression of a signal vector, \mathbf{x} , by identical scalar quantization of its components, x_i :

$$LMSE^{SQV}(\mathbf{A}, \mathbf{b}; \bar{\eta}) \triangleq$$

$$= \frac{1}{|\bar{\eta}|} \int_{\bar{\eta}} \|Q(\mathbf{x}) - \mathbf{A}\mathbf{x} - \mathbf{b}\|_{2}^{2} d\mathbf{x}$$

$$= \frac{1}{|\bar{\eta}|} \sum_{i=1}^{N} \int_{\bar{\eta}} \left(q(x_{i}) - \mathbf{a}_{i}^{T}\mathbf{x} - b_{i}\right)^{2} d\mathbf{x} \qquad (43)$$

where $\bar{\eta}$ is an arbitrary shaped approximation area, and the last equality relies on the separability of $Q(\cdot)$ and the L_2 -norm definition.

Equations (41) and (43) clearly show that the MSE of approximating transform-coding (41) reduces to the linearization-error of scalar quantization of the transform coefficients, namely,

$$LMSE^{TC}(\mathbf{A}, \mathbf{b}; \eta(\mathbf{x}_{0}, \delta))$$

$$= LMSE^{SQV}(\hat{\mathbf{A}}, \hat{\mathbf{b}}; \hat{\eta}(\mathbf{x}_{0}, \delta))$$

$$= \frac{1}{|\hat{\eta}(\mathbf{x}_{0}, \delta)|} \sum_{i=1}^{N} \int_{\hat{\eta}(\mathbf{x}_{0}, \delta)} (q(\hat{x}_{i}) - \hat{\mathbf{a}}_{i}^{T} \hat{\mathbf{x}} - \hat{b}_{i})^{2} d\hat{\mathbf{x}} \quad (44)$$

where $\hat{\mathbf{a}}_{i}^{T}$ is the i^{th} row of $\hat{\mathbf{A}}$, and \hat{b}_{i} is the i^{th} element of the vector $\hat{\mathbf{b}}$.

While the separability of $Q(\cdot)$ was utilized to have integrals in (44) that consider quantization of single transformcoefficients, the integration is still over a multidimensional area that is not necessarily separable (i.e., not aligned with the axes). We can remedy this by starting from an appropriately rotated area in the signal-domain, $\eta_U(\mathbf{x}_0, \delta)$, such that its transform-domain counterpart is aligned with the axes (see Fig. 4):

$$\hat{\eta}_U(\mathbf{x}_0, \delta) = \left\{ \hat{\mathbf{x}} \mid \| \hat{\mathbf{x}} - \hat{\mathbf{x}}_0 \|_{\infty} \le \delta \right\}.$$
(45)

Note that $\eta_U(\mathbf{x}_0, \delta)$ is not necessarily the optimally shaped approximation area as it is used here for the analytic simplicity of having full separability in the transform domain.



Fig. 4. Transformation of the approximation area. Exemplified in \mathbb{R}^2 for the unitary transform of 45°-rotation and a signal-domain area that is rotated in accordance to the transform.

We continue our transform-domain analysis by adopting this separable integration-area.

Recall that we look for the optimal linear approximation of the signal-domain function $C(\mathbf{x})$. This is obtainable by finding the optimal transform-domain parameters $\hat{\mathbf{A}}^*$ and $\hat{\mathbf{b}}^*$ and then transforming them back to the signal domain. Following this strategy we first pose the componentwise optimality demands in the transform domain:

$$\frac{\partial}{\partial \hat{a}_{ij}} LMSE^{SQV} \left(\hat{\mathbf{A}}, \hat{\mathbf{b}}; \hat{\eta}_U(\mathbf{x}_0, \delta) \right) = 0 \quad \text{for } i, j = 1, ..., N$$
$$\frac{\partial}{\partial \hat{b}_i} LMSE^{SQV} \left(\hat{\mathbf{A}}, \hat{\mathbf{b}}; \hat{\eta}_U(\mathbf{x}_0, \delta) \right) = 0 \quad \text{for } i = 1, ..., N$$
(46)

Some calculations show that the solution satisfying the optimality conditions consists of a diagonal matrix $\hat{\mathbf{A}}^*$ (i.e., $\hat{a}_{ij}^* = 0$ for $i \neq j$) such that the parameter pair $(\hat{a}_{ii}^*, \hat{b}_i^*)$ is the one obtained for optimal approximation of a single-variable quantizer over the interval $[\hat{\mathbf{x}}_0^{(i)} - \delta, \hat{\mathbf{x}}_0^{(i)} + \delta]$ as generally given in (15)-(16). Then, the signal-domain parameters are given as

$$\mathbf{A}^* = \mathbf{U}\hat{\mathbf{A}}^*\mathbf{U}^T = \sum_{i=1}^N \hat{a}_{ii}^*\mathbf{u}_i\mathbf{u}_i^T$$
(47)

$$\mathbf{b}^* = \mathbf{U}\hat{\mathbf{b}}^*. \tag{48}$$

where the last equality in (47) is due to the diagonality of $\hat{\mathbf{A}}^*$ and \mathbf{u}_i denotes the i^{th} column of U. The corresponding optimal error is equivalent in the signal and transform domains, hence can be expressed in a simplified form as

T C

$$LMSE^{IC} \left(\mathbf{A}^{*}, \mathbf{b}^{*}; \eta_{U}(\mathbf{x}_{0}, \delta) \right)$$

= $LMSE^{SQV} \left(\hat{\mathbf{A}}^{*}, \hat{\mathbf{b}}^{*}; \hat{\eta}_{U}(\mathbf{x}_{0}, \delta) \right)$
= $\frac{1}{2\delta} \sum_{i=1}^{N} \int_{\hat{\mathbf{x}}_{0}^{(i)} - \delta}^{\hat{\mathbf{x}}_{0}^{(i)} + \delta} \left(q(\hat{x}_{i}) - \hat{a}_{ii}^{*} \hat{x}_{i} - \hat{b}_{i}^{*} \right)^{2} d\hat{\mathbf{x}}$
= $\sum_{i=1}^{N} LMSE^{SQ} \left(\hat{a}_{ii}^{*}, \hat{b}_{i}^{*}; \eta(\hat{\mathbf{x}}_{0}^{(i)}, \delta) \right)$ (49)

The last expression exhibits the approximation error of transform-coding as the sum of the errors of the separate linearization of the scalar quantization of the transform-domain coefficients. Although the assumed scenario includes equal quantization procedure for all the coefficients, the contributed



Fig. 5. Overall MSE of optimal linearization of the exemplary transform coding procedure (for signals in \mathbb{R}^2 and equal transform-domain quantizers).

errors by the various elements are different as each has its own scalar approximation-point $\hat{\mathbf{x}}_{0}^{(i)}$ located differently with respect to the quantization lattice.

Let us exemplify the latter analysis on a transform coder of two-component signals (i.e., $\mathbf{x} \in \mathbb{R}^2$), that scalary applies the normalized two-level quantizer that was studied above (see Eq. (19)) on the two components in the domain of the 45° rotation matrix that takes the 2×2 form of $U_{\pi/4} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$. The approximation is around $\mathbf{x}_0 = \mathbf{U}_{\pi/4} \hat{\mathbf{x}}_0$ in a 45° -rotated square neighborhood defined by the $\eta_U(\mathbf{x}_0, \delta)$ (see Fig. 4). We further define the first component of \hat{x}_0 to vary and fix the second on the value of 15, i.e., $\hat{\mathbf{x}}_0 = \begin{bmatrix} \hat{\mathbf{x}}_0^{(1)} \\ 15 \end{bmatrix}$. The overall linearization error as function of δ and $\hat{\mathbf{x}}_0^{(1)}$ (Fig. 5) shows that it combines the errors of the scalar linearization of the transform coefficients (Figs. 6e-6f). The corresponding parameters in the signal domain (where the matrix A is not necessarily diagonal) are given in Fig. 7. Again, the results generalize the previous observations by demonstrating that minimal MSE is obtained for approximation over the minimal area that includes the nearest non-differentiable point of the compression function (see Fig. 5).

We now further develop the discussed transform-coder to the common procedure where the transform-coefficients are uniformly quantized according to different step sizes, $\{\Delta_i\}_{i=1}^N$, getting coarser for higher frequencies, i.e., $\Delta_i \leq \Delta_i$ for i < j. Let us consider N-length signal vectors and analyze the linear approximation over the $\eta_U(\mathbf{x}_0, \delta)$ neighborhood, where A^* is diagonal. Accordingly, in this case, quantization of each transform-coefficient is linearized separately over a one-dimensional interval of size 2δ . However, due to different quantization-steps and approximation-points, the $\{\hat{a}_{ii}^*\}_{i=1}^N$ values vary. We simplify the discussion by approximating around a vector \mathbf{x}_0 with components residing near the middle of the scalar decision-regions of the transform-domain quantizers. Then, relying on the above analysis of the one-dimensional uniform quantizer, we note the following behavior of the sequence $\{\hat{a}_{ii}^*\}_{i=1}^N$ for $2\delta \approx \Delta_K$ (K > 1). First, for some integer L $(1 \leq L < K)$, the relation $\Delta_i \ll \Delta_K \approx 2\delta$



Fig. 6. Transform domain parameters of the optimal linear approximation of the exemplary transform coding procedure (for signals in \mathbb{R}^2 and equal transform-domain quantizers). (a)-(b) describe the diagonal elements of the 2×2 matrix $\hat{\mathbf{A}}$, and (c)-(d) show the values of $\hat{\mathbf{b}}$'s components. (e)-(f) show the corresponding approximation errors of the twotransform domain elements. (a) \hat{a}_{11}^* . (b) \hat{a}_{22}^* . (c) \hat{b}_{1}^* . (d) \hat{b}_{2}^* . (e) $LMSE^{SQ}(\hat{a}_{11}^*, \hat{b}^*; \eta(\hat{\mathbf{x}}_{0}^{(1)}, \delta))$. (f) $LMSE^{SQ}(\hat{a}_{22}^*, \hat{b}_{22}^*; \eta(\hat{\mathbf{x}}_{01}^{(2)}, \delta))$.

holds for any $i \leq L$, and therefore, the optimal parameters are $\hat{a}_{ii}^* \approx 1$ and $\hat{b}_i^* \approx 0$. Second, for $L < i \leq K$, the 2δ value is still greater than Δ_i , however relatively closer, hence, the corresponding \hat{a}_{ii}^* values fluctuate. Finally, the i > K coefficients have quantization steps that are greater than 2δ , and accordingly, $\hat{a}_{ii}^* \approx 0$. The latter qualitative analysis lets us to interpret the local-approximation of the transform-coder, \hat{A}^* , as a low-pass filter that depends on δ and the approximation point. Furthermore, the numerical results (Fig. 8) demonstrate the above by showing preservation of low frequencies, an unstable transition phase, and attenuation of high-frequency components. Note that for too low or too high values of δ the filter has a all-stop (Fig. 8a) or all-pass (Fig. 8d) behavior, respectively.

We conclude by considering the signal-domain filter \mathbf{A}^* related to $\hat{\mathbf{A}}^*$ by the inverse-transformation in (47). When the compression utilizes the Discrete Fourier Transform (DFT) and the approximation is over $\eta_U(\mathbf{x}_0, \delta)$, then the diagonal matrix $\hat{\mathbf{A}}^*$ yields a circulant \mathbf{A}^* . While the latter involves complex-valued calculations, coding using Discrete Cosine Transform (DCT) keeps the procedure over the reals. The signal-domain filter, \mathbf{A}^* , of the DCT-based coding is



Fig. 7. Signal domain parameters of the optimal linear approximation of the exemplary transform coding procedure (for signals in \mathbb{R}^2 and equal transformdomain quantizers). (a)-(d) describe the components of the 2x2 matrix **A**, and (e)-(f) show the values of **b**'s components. (a) a_{11}^* . (b) a_{12}^* . (c) a_{21}^* . (d) a_{22}^* . (e) b_1^* . (f) b_2^* .

exemplified in Fig. 9 showing an approximately Toeplitz structure.

IV. EXPERIMENTAL RESULTS

In this section we demonstrate the performance of the proposed postprocessing method by presenting results obtained in conjunction with various compression methods. We start by considering the simplistic compression procedures of scalar quantization and one-dimensional transform coding. Then, we proceed to the leading image compression standards: JPEG, JPEG2000 and the recent HEVC.

In all experiments we use the BM3D method [27] as the denoiser. Since the proposed technique uses a well established denoiser as a subroutine, we compare our method with a single application of this denoiser as a postprocessing procedure. This approach is further strengthened by endorsing the denoiser with an oracle capability by searching for the best parameter in terms of maximal PSNR result. More specifically, this oracle denoiser optimizes its output PSNR based on the knowledge of the precompressed image, a capability that cannot be applied in a real postprocessing task.

The computational complexity of our method is mainly determined by the complexity levels of the utilized denoiser



Fig. 8. Interpretation of a diagonal $\hat{\mathbf{A}}^*$ as a transform-domain filter that depends on δ . Here N = 32 and the i^{th} quantization step is $\Delta_i = 2^{i/4}$. (a) $\delta = 0.5$. (b) $\delta = 5$. (c) $\delta = 50$. (d) $\delta = 500$.



Fig. 9. Interpretation of A^* as a signal-domain filter that depends on δ . Presented here for the case of $\delta = 50$ from Fig. 8c, incorporated in a DCT-based coding.

and the Jacobian estimation procedure. The latter further depends on the implementation of the compressiondecompression method, as it is repeatedly applied according to (9). In addition, equation (9) exhibits also the effect of the size of the set S_{δ} utilized for approximating a single column of the Jacobian. Since the number of Jacobian columns is as the number of signal samples (denoted as N), a straightforward computation of the Jacobian is costly and requires N calculations of (9). Furthermore, the Jacobian matrix is of $N \times N$ size, and is often too large to allow accurate solution of (10). Fortunately, the computational requirements of the estimation of the entire Jacobian matrix can be relaxed for many compression methods that operate independently on adjacent blocks. Specifically, the Jacobian becomes a blockdiagonal matrix and, therefore, its columns can be arranged in independent subsets for concurrent computation. This reduces the number of compression-decompression applications to the order of the block size. Moreover, the block-diagonal

 TABLE I

 Experimental Settings for the Examined Compression Methods

Compression Method	Affecting Factors	Max. Iterations	$S_{\delta} = \left\{ 0.1 \tilde{\Delta} k \right\}_{k=1}^{5}$	λ	β	μ	$B_H \times B_W$
Scalar Quantization	r bit-rate (bpp) Δ quantizer step size	6	$\tilde{\Delta} = \Delta$	0.01	$500 \cdot 2^{-2r}$	$5 \cdot 10^{-4} \cdot 2^{0.6r}$	1×1
Simplistic	Δ quantizer step size (in transform domain)				Aligned approx. area: $200 \cdot 2^{-0.5\tilde{r}}$		
Transform Coding	effective bit-rate $\tilde{r} = 16 - log_2(\Delta)$	10	$\tilde{\Delta} = \Delta$	0.03	Rotated approx. area: $500 \cdot 2^{-0.5\tilde{r}}$	$5 \cdot 10^{-5} \cdot 2^{0.8\tilde{r}}$	2×1
JPEG	r bit-rate (bpp)	8	$\tilde{\Delta} = 135/r$	0.15	$2 \cdot r^{-1}$	$0.01 \cdot 2^r$	8×8
JPEG2000	r bit-rate (bpp)	8	$\tilde{\Delta} = 100/r$	0.15	$5 \cdot r^{-1}$	$0.3 \cdot 2^r$	8×8
HEVC	r bit-rate (bpp)	8	$\tilde{\Delta} = 50/r$	0.15	$5 \cdot r^{-1}$	$0.3 \cdot 2^r$	64×64

structure of the Jacobian allows to decompose the computation of (10) to handle each block separately. Furthermore, this block-diagonal structure can be assumed even for compression methods that do not conform with it (e.g., JPEG2000), and thus somewhat compromising the postprocessing result, in order to offer a reasonable run-time. The (possibly assumed) block size of the compression procedure is denoted here as $B_H \times B_W$, and yields a Jacobian with blocks of size $B_H B_W \times B_H B_W$ along its main diagonal.

The code was implemented in Matlab. While the settings differ for the various compression methods, a similar stopping criterion is applied. In (3) we introduced the scaled dual-variable of the i^{th} iteration, $\mathbf{u}_i \in \mathbb{R}^n$. We here denote $\Delta \mathbf{u}_i = \frac{1}{N} \|\mathbf{u}_i - \mathbf{u}_{i-1}\|_1$ and set the algorithm termination conditions to be at one of the following: $\Delta \mathbf{u}_i < 0.05$, $\Delta \mathbf{u}_i > \Delta \mathbf{u}_{i-1}$ or some maximal number of iterations attained.

The remaining parameters are set for each compression method as specified in Table I. While the relation between the parameters to the compression method is complex, one can claim that the parameters express the non-differentiable nature of the compression function. For example, HEVC compression, which is an intricate compression method, needs smaller δ values in the Jacobian approximation and a higher μ , both constraining the linearization to be more local than for the other simpler compression methods. Furthermore, the parameter settings consider the compression bit-rate, as this quantity reflects the complexity of the given image with respect to the specific compression procedure. Accordingly, the formulas in Table I were empirically determined to provide an adequate performance.

A. Simplistic Compression Procedures

1) Scalar Quantization: We begin with the elementary compression procedure of applying uniform scalar quantization (as formulated in (27)) on the signal samples. Motivated by the analysis in section III-C, we define here the approximation interval to be no longer than the quantization step Δ , as considering only the nearest non-differentiable point yields a useful linearization. Accordingly, the derivative (which is scalar here) is approximated using (9) and $S_{\delta} = \{0.1 \Delta k\}_{k=1}^{5}$. Our technique achieved impressive PSNR improvements (Table II) over the entire bit-rate range, and consistently passed the oracle denoiser. Visually, the false-contouring artifacts were significantly reduced (Fig. 10).

TABLE II SIMPLISTIC EXPERIMENT: PSNR COMPARISON FOR SCALAR QUANTIZATION

Image 256x256	Bit-Rate	No Postprocessing	Oracle Denoiser	Proposed Method
Lena	2	22.86	24.54	24.84
	3	28.90	30.80	31.02
	4	34.64	37.19	37.41
	5	40.72	42.82	43.02
Barbara	2	23.47	25.66	25.93
	3	28.48	30.94	31.11
	4	34.72	37.29	37.38
	5	40.74	42.55	42.64



Fig. 10. Reconstruction of Lena (256x256) from scalar quantization at 4bpp. (a) Scalar Quantization (34.64dB). (b) Postprocessing Result (37.41dB).

2) One-Dimensional Transform Coding: Now we extend the examined compression method by performing scalar quantization in the transform domain. Specifically, we split the image into nonoverlapping one-dimensional vertical vectors of 2 pixels. Then we compress them separately by transform coding them using $\mathbf{U}_{\pi/4} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$, followed by applying uniform quantization with identical step size to all the coefficients. We evaluated here the algorithm performance for the two types of approximation area that were discussed in section III-D: a square area aligned with the axes (i.e., $\eta(\mathbf{x}_0, \delta)$), and a 45°-rotated squared area (i.e., $\hat{\eta}_U(\mathbf{x}_0, \delta)$) that allows to calculate the linearization in the transform domain and then transforming the parameters back to the signal domain using (47)-(48). Both options used square areas of the same size by setting $S_{\delta} = \{0.1 \Delta k\}_{k=1}^{5}$, where Δ is the quantizer step size in the transform domain. The two

TABLE III
SIMPLISTIC EXPERIMENT: PSNR COMPARISON FOR TRANSFORM CODING OF TWO-COMPONENT VECTORS

Image 256x256	Quantizer Step Δ	No Postprocessing	Oracle Denoiser	Proposed Method Aligned Approx. Area	Proposed Method Rotated Approx. Area
Lena	20	31.22	35.83	37.00	36.78
	30	27.45	32.21	33.37	33.00
	40	24.88	29.92	30.92	30.29
Barbara	15	34.50	38.00	38.25	38.12
	20	31.79	35.65	36.42	36.31
	30	28.05	32.33	33.33	33.28

TABLE IV JPEG: Result Comparison

Image		JPEG		Oracle Denoiser		Foi et al. [17]		Zhang et al. [18]		Proposed Method	
512x512	Bit-Rate	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Lena	0.173	27.33	0.7367	28.77	0.7887	28.95	0.8040	29.07	0.8054	28.90	0.7965
	0.245	30.41	0.8183	31.82	0.8555	31.84	0.8591	31.97	0.8611	31.63	0.8518
	0.363	32.96	0.8735	34.12	0.8932	34.06	0.8927	34.24	0.8953	34.32	0.8958
	0.511	34.76	0.9036	35.64	0.9130	35.55	0.9124	35.82	0.9155	35.87	0.9160
	0.638	35.81	0.9188	36.52	0.9238	36.44	0.9235	36.77	0.9270	36.81	0.9271
	0.807	36.86	0.9314	37.42	0.9334	37.34	0.9332	37.73	0.9374	37.79	0.9373
	1.157	38.54	0.9476	38.91	0.9474	38.82	0.9466	39.25	0.9511	39.23	0.9500
Barbara	0.227	23.86	0.6642	25.40	0.7142	24.97	0.7212	25.33	0.7290	24.50	0.7008
	0.338	25.70	0.7710	27.26	0.8038	26.54	0.8076	27.32	0.8161	26.40	0.8063
	0.537	28.25	0.8559	30.03	0.8761	28.91	0.8774	30.19	0.8859	29.67	0.8904
	0.764	30.89	0.9060	32.49	0.9175	31.42	0.9178	32.80	0.9259	32.59	0.9286
	0.938	32.54	0.9273	33.94	0.9338	33.02	0.9350	34.32	0.9419	34.33	0.9439
	1.149	34.22	0.9442	35.34	0.9464	34.64	0.9491	35.81	0.9545	35.94	0.9559
	1.552	36.88	0.9625	37.65	0.9619	37.22	0.9650	38.18	0.9681	38.31	0.9681
Boat	0.187	25.56	0.6563	26.70	0.6962	26.75	0.7013	26.83	0.7048	26.74	0.7015
	0.291	28.13	0.7580	29.19	0.7883	29.14	0.7867	29.26	0.7921	29.03	0.7883
	0.460	30.49	0.8301	31.37	0.8471	31.29	0.8447	31.52	0.8501	31.61	0.8550
	0.663	32.35	0.8691	33.14	0.8785	33.04	0.8779	33.35	0.8823	33.47	0.8854
	0.825	33.50	0.8880	34.16	0.8926	34.07	0.8938	34.38	0.8975	34.51	0.8992
	1.035	34.63	0.9042	35.19	0.9062	35.10	0.9071	35.40	0.9109	35.50	0.9109
	1.497	36.43	0.9289	36.91	0.9292	36.82	0.9311	37.10	0.9338	37.07	0.9309
Bridge	0.231	23.06	0.5713	23.76	0.5848	23.79	0.5807	23.85	0.5926	23.60	0.5921
	0.409	25.13	0.7108	25.71	0.7175	25.72	0.7157	25.76	0.7234	25.62	0.7284
	0.688	27.01	0.8093	27.48	0.8118	27.50	0.8157	27.53	0.8179	27.53	0.8210
	1.003	28.50	0.8634	28.88	0.8603	28.88	0.8692	29.00	0.8702	29.02	0.8726
	1.261	29.54	0.8917	29.95	0.8893	29.89	0.8976	30.11	0.8989	29.92	0.8999
	1.579	30.80	0.9166	31.23	0.9153	31.09	0.9215	31.41	0.9233	31.17	0.9233
	2.197	33.32	0.9490	33.78	0.9474	33.57	0.9522	34.00	0.9544	33.72	0.9540
Peppers	0.176	27.17	0.7078	28.75	0.7748	29.04	0.7906	29.17	0.7890	28.86	0.7802
	0.246	30.14	0.7839	31.50	0.8273	31.69	0.8322	31.77	0.8316	31.68	0.8316
	0.361	32.44	0.8354	33.45	0.8564	33.52	0.8561	33.67	0.8590	33.63	0.8584
	0.513	33.92	0.8651	34.61	0.8730	34.64	0.8722	34.83	0.8775	34.81	0.8763
	0.650	34.77	0.8809	35.30	0.8829	35.31	0.8824	35.55	0.8892	35.55	0.8878
	0.832	35.59	0.8949	35.96	0.8941	35.94	0.8927	36.19	0.9001	36.17	0.8980
	1.203	36.79	0.9135	37.03	0.9111	37.01	0.9105	37.23	0.9161	37.13	0.9130

area types achieved better results than the oracle denoiser. In that sense, our approach is somewhat robust to the area shape (however, not necessarily to its size). In addition, employing area that is aligned with the axes in the signal domain consistently obtained higher PSNR than the rotated area (Table III). The latter observation will motivate us to use aligned-cubic approximation areas also for more complex compression techniques that will follow next.



Fig. 11. Reconstruction of Lena (512x512) from JPEG compression at 0.363bpp. (a) JPEG (32.96dB). (b) Postprocessing Result (34.32dB).

B. JPEG

This well known standard [28] is a relatively straightforward implementation of a two-dimensional transform coding on 8×8 blocks of the image. Specifically, the quantization is performed in the DCT domain where each coefficient has its own quantization step. As the JPEG extends the oversimplified procedure in subsection IV-A2, our postprocessing method is expected to provide good results here. Indeed, the experiments show the impressive gains of the suggested method that compete with the prominent techniques from [17], [18] (see Table IV). The comparison to [17] and [18] indirectly considers additional methods such as [1], [3], [14], and [37] that were already surpassed by [17] and/or [18]. Moreover, while many competitive methods (e.g., [1], [3], [14], and [17]) are mainly intended to low bit-rate compression, our method handles the entire bit-rate range and excels for medium and high bit-rates (Table IV). The thorough evaluation here is based on PSNR values, as well as on the perceptual metric of Structural Similarity (SSIM) [38].

Since JPEG applies transform coding on non-overlapping 8×8 blocks, its Jacobian matrix is indeed block diagonal. In addition, the sufficiently small blocks provide a computationally efficient structure that does not need to be simplified further. Consequently, the run-time of the JPEG postprocessing (Matlab implementation) is rather reasonable and is usually about 1-2 minutes for a 512×512 image.

C. JPEG2000

This efficient standard [29] applies transform coding in the wavelet domain for relatively large signal blocks (also known as tiles) of at least 128×128 size. Not only the tile size affects the compression run-time, it also impairs the suggested parallelism optimization, as it is beneficial for small block sizes. Nevertheless, it is still recommended to reduce the computational cost by concurrent computation of the Jacobian columns in relatively large subgroups that inevitably contain dependent elements. Our experiments included postprocessing of images compressed using JPEG2000 compression (via the Kakadu software [39]) without any tiling. However, the Jacobian was estimated by assuming independent 8×8 blocks, where this reduced accuracy yielded considerable relief in the computational burden (the postprocessing took 5-8 minutes for

 TABLE V

 JPEG2000: Comparison of PSNR Gains at Low Bit-Rates

Image 512x512 ¹	Bit Rate ²	JPEG2000 ³	Zhai et al. [19]	Kwon et al. [20]	Proposed Method
Lena	0.10	29.86	0.03	0.50	0.40
	0.15	31.63	-0.34	0.54	0.39
	0.20	33.01	0.08	NA	0.47
Peppers	0.10	29.58	0.31	0.47	0.47
	0.15	31.33	0.04	0.62	0.41
	0.20	32.51	-0.27	NA	0.46
Bridge	0.10	22.81	-0.05	0.04	0.06
_	0.15	23.76	-0.13	0.10	0.09
	0.20	24.34	-0.17	NA	0.11

¹ The results presented in [19], [20] for JPEG2000 postprocessing are for somewhat different versions of the widely-used Lena, Peppers and Bridge images. Accordingly, in this table, and only here, we refer to these images that were specified in [20].

² The bit-rate here is the input given to the Kakadu software, and is not necessarily the accurate output bit-rate.

³ As explained in [20], the PSNR values of the compression using the Kakadu software are slightly different in the various papers. Accordingly, the comparison is for the PSNR gains.

a 512×512 image). The reconstruction PSNR of our method reached up to 0.7dB improvement of the JPEG2000 output (e.g., see Fig. 12). The results in [19] and [20] were provided to postprocessing of low bit-rate compression. Therefore, we first compare our results to these from [19] and [20] (Table V) according to their experimental settings, and then show results for higher bit-rates where our method is even more effective (Table VI). Table V exhibits that our method outperforms [19] and competitive with the technique from [20]. In addition, our results for higher bit-rates (Table VI) compete with the oracle denoiser. These results are encouraging since the oracle denoiser needs the precompressed image and, therefore, is not suitable for the common compression applications. Furthermore, the results in Tables V and VI establish our technique as suitable for a wide range of bit-rates. The restoration results visually demonstrated the artifact reduction using our method, specifically, handling of the ringing artifact (Fig. 12).

D. HEVC

This state-of-the-art coding standard offers a profile of stillimage compression [30], [31]. The HEVC applies spatial hybrid-coding on the image by combining a rich prediction capability with transform coding of the prediction residuals. In addition, the image is divided into large blocks (also known as coding units) that are further recursively partitioned into rectangular blocks in various sizes. Therefore, our Jacobian estimation is set to work on independent blocks of size 64×64 , and thus the corresponding run-time was higher than for the previous compression methods. More specifically, postprocessing an HEVC-compressed image took several hours, in contrast to few minutes as needed for the previous compression methods. Accordingly, we stress that our purpose here is to demonstrate the conceptual suitability of our method to compression techniques that are significantly more intricate than transform coding.

The results here are for HEVC-compression using the software library in [40]. Again, our postprocessing results

 TABLE VI

 JPEG2000: Result Comparison at Medium/High Bit-Rates

-		THE INCOME			acle	Proposed		
Image	Bit	JPEO	G2000	Den	oiser	Me	thod	
512x512	Kate	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	
Lena	0.30	34.87	0.8983	35.20	0.9005	35.24	0.9019	
	0.40	36.12	0.9142	36.43	0.9160	36.42	0.9161	
	0.50	37.26	0.9255	37.49	0.9262	37.39	0.9254	
	0.60	37.97	0.9326	38.16	0.9331	37.99	0.9313	
	0.70	38.69	0.9392	38.87	0.9398	38.58	0.9370	
Barbara	0.30	29.16	0.8506	30.00	0.8607	29.76	0.8635	
	0.40	30.79	0.8791	31.70	0.8897	31.51	0.8915	
	0.50	32.16	0.9049	33.09	0.9131	32.81	0.9134	
	0.60	33.30	0.9199	34.20	0.9263	34.03	0.9277	
	0.70	34.37	0.9296	35.17	0.9342	35.05	0.9358	
Boat	0.30	30.87	0.8204	31.22	0.8237	31.25	0.8272	
	0.40	32.29	0.8508	32.61	0.8535	32.62	0.8557	
	0.50	33.32	0.8710	33.58	0.8729	33.63	0.8750	
	0.60	34.17	0.8858	34.44	0.8864	34.48	0.8889	
	0.70	34.90	0.8975	35.13	0.8973	35.15	0.8993	
Bridge	0.30	25.42	0.6922	25.50	0.6848	25.52	0.6920	
	0.40	26.36	0.7510	26.41	0.7397	26.47	0.7509	
	0.50	27.24	0.7852	27.35	0.7768	27.39	0.7865	
	0.60	27.89	0.8143	28.02	0.8077	28.04	0.8160	
	0.70	28.50	0.8391	28.64	0.8337	28.65	0.8408	
Peppers	0.30	34.11	0.8587	34.37	0.8624	34.39	0.8625	
	0.40	35.05	0.8722	35.26	0.8739	35.24	0.8737	
	0.50	35.80	0.8833	35.96	0.8838	35.90	0.8832	
	0.60	36.34	0.8955	36.44	0.8936	36.42	0.8947	
	0.70	36.83	0.9056	36.91	0.9027	36.93	0.9052	

TABLE VII HEVC: Result Comparison

Bit	Bit HEVC		Or Den	acle oiser	Proposed Method		
Rate	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	
0.177	25.79	0.7278	25.91	0.7331	25.89	0.7337	
0.340	28.91	0.8357	28.97	0.8349	29.00	0.8359	
0.639	32.71	0.9180	32.84	0.9181	32.92	0.9202	
1.046	36.31	0.9577	36.44	0.9572	36.51	0.9583	
0.120	27.55	0.7658	27.71	0.7726	27.76	0.7749	
0.206	30.22	0.8420	30.39	0.8471	30.45	0.8466	
0.401	33.30	0.9163	33.48	0.9181	33.54	0.9187	
0.746	36.81	0.9554	37.04	0.9570	37.12	0.9581	
0.212	25.29	0.7566	25.33	0.7568	25.35	0.7574	
0.416	28.54	0.8640	28.59	0.8614	28.67	0.8650	
0.735	32.44	0.9346	32.55	0.9352	32.68	0.9369	
1.186	36.51	0.9682	36.61	0.9684	36.68	0.9692	
0.198	25.08	0.7172	25.11	0.7142	25.10	0.7148	
0.393	28.10	0.8287	28.10	0.8287	28.11	0.8268	
0.746	31.64	0.9151	31.64	0.9151	31.71	0.9144	
1.289	35.68	0.9653	35.69	0.9653	35.71	0.9649	
0.1323	28.40	0.8328	28.53	0.8409	28.61	0.8408	
0.2285	31.51	0.8947	31.71	0.8993	31.81	0.8987	
0.3813	34.73	0.9330	34.92	0.9354	34.89	0.9340	
0.6284	38.01	0.9563	38.18	0.9569	38.10	0.9560	
	Bit Rate 0.177 0.340 0.639 1.046 0.120 0.206 0.401 0.746 0.212 0.416 0.735 1.186 0.198 0.393 0.746 1.289 0.1323 0.2285 0.3813 0.6284	Bit Rate HI PSNR 0.177 25.79 0.340 28.91 0.639 32.71 1.046 36.31 0.120 27.55 0.206 30.22 0.401 33.30 0.746 36.81 0.212 25.29 0.416 28.54 0.735 32.44 1.186 36.51 0.198 25.08 0.393 28.10 0.746 31.64 1.289 35.68 0.1323 28.40 0.2285 31.51 0.3813 34.73 0.6284 38.01	Bit Rate HEVC PSNR SSIM 0.177 25.79 0.7278 0.340 28.91 0.8357 0.639 32.71 0.9180 1.046 36.31 0.9577 0.120 27.55 0.7658 0.206 30.22 0.8420 0.401 33.30 0.9163 0.746 36.81 0.9554 0.212 25.29 0.7566 0.416 28.54 0.8640 0.735 32.44 0.9346 1.186 36.51 0.9682 0.198 25.08 0.7172 0.393 28.10 0.8287 0.746 31.64 0.9151 1.289 35.68 0.9653 0.1323 28.40 0.8328 0.2285 31.51 0.8947 0.3813 34.73 0.9330 0.6284 38.01 0.9553	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	Bit Rate HEVC Oracle Denoiser PSNR SSIM PSNR SSIM 0.177 25.79 0.7278 25.91 0.7331 0.340 28.91 0.8357 28.97 0.8349 0.639 32.71 0.9180 32.84 0.9181 1.046 36.31 0.9577 36.44 0.9572 0.120 27.55 0.7658 27.71 0.7726 0.206 30.22 0.8420 30.39 0.8471 0.401 33.30 0.9163 33.48 0.9181 0.746 36.81 0.9554 37.04 0.9570 0.212 25.29 0.7566 25.33 0.7568 0.416 28.54 0.8640 28.59 0.8614 0.735 32.44 0.9346 32.55 0.9352 1.186 36.51 0.9682 36.61 0.9684 0.198 25.08 0.7172 25.11 0.7142 0.393 28.10	Bit Rate $OracleDenoiser PropMet BitRate HEVC Denoiser Met PSNR SSIM PSNR SSIM PSNR SSIM 0.177 25.79 0.7278 25.91 0.7331 25.89 0.340 28.91 0.8357 28.97 0.8349 29.00 0.639 32.71 0.9180 32.84 0.9181 32.92 1.046 36.31 0.9577 36.44 0.9572 36.51 0.120 27.55 0.7658 27.71 0.7726 27.76 0.206 30.22 0.8420 30.39 0.8471 30.45 0.401 33.30 0.9163 33.48 0.9181 33.54 0.746 36.81 0.9554 37.04 0.9570 37.12 0.212 25.29 0.7566 25.33 0.7568 25.35 0.416 28.54 0.8640 28.59 0.8614 28.67 0.735 32.44 $	

Here the images are the 128×128 portions taken from the center of the 256×256 images.



Fig. 12. Reconstruction of Barbara (512x512) from JPEG2000 compression at 0.40bpp. (a) JPEG2000 (30.79dB). (b) Postprocessing Result (31.51dB). (c) JPEG2000 (Zoomed-in). (d) Postprocessing Result (Zoomed-in).

reached up to 0.3dB gain in PSNR and often exceeded the oracle denoiser, as shown in the PSNR and SSIM comparison in Table VII. Figure 13 visually demonstrates our method's



Fig. 13. Reconstruction of Lena (128x128 Portion) from HEVC compression at 0.639bpp. (a) HEVC (32.71dB). (b) Postprocessing Result (32.92dB).

treatment of the delicate artifacts of the HEVC. To the best of our knowledge, no other artifact-reduction techniques for the HEVC still-image profile have been proposed yet, as it is a recent standard.

To summarize this section, the extensive experiments established the proposed compression-artifact reduction technique as a generic method that achieves cutting-edge results for any relevant image compression and over the entire bit-rate range.

V. CONCLUSION

In this paper we proposed a novel postprocessing method for reducing artifacts in compressed images. The task was formulated as a regularized inverse problem, that was subsequently transformed into an iterative form by relying on the ADMM and the Plug-and-Play frameworks. The resulting generic algorithm separately treats the inversion and the regularization, where the latter is implemented by sequentially applying an existing state-of-the-art Gaussian denoiser. For practicality we simplified the inversion step by representing the nonlinear compression-decompression procedure using a linear approximation. Furthermore, we provided a comprehensive mathematical analysis for linear approximation of simplified quantization and transform-coding operations. We demonstrated our approach for image compression and presented experimental-results showing impressive gains, that improve upon state-of-the-art postprocessing results for leading image compression standards.

REFERENCES

- A. Z. Averbuch, A. Schclar, and D. L. Donoho, "Deblocking of block-transform compressed images using weighted sums of symmetrically aligned pixels," *IEEE Trans. Image Process.*, vol. 14, no. 2, pp. 200–212, Feb. 2005.
 Y.-C. Liaw, W. Lo, and J. Z. C. Lai, "Image restoration of compressed
- [2] Y.-C. Liaw, W. Lo, and J. Z. C. Lai, "Image restoration of compressed image using classified vector quantization," *Pattern Recognit.*, vol. 35, no. 2, pp. 329–340, Feb. 2002.
- [3] T. Chen, H. R. Wu, and B. Qiu, "Adaptive postfiltering of transform coefficients for the reduction of blocking artifacts," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 11, no. 5, pp. 594–602, May 2001.
- [4] K. Lee, D. S. Kim, and T. Kim, "A new postprocessing algorithm based on regression functions," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, vol. 4. May 2002, pp. IV-3684–IV-3687.
- [5] G. A. Triantaffilidis, D. Sampson, D. Tzovaras, and M. G. Strintzis, "Blockiness reduction in JPEG coded images," in *Proc. 14th Int. Conf. Digit. Signal Process.*, vol. 2, 2002, pp. 1325–1328.
- [6] F. Alter, S. Durand, and J. Froment, "Adapted total variation for artifact free decompression of JPEG images," J. Math. Imag. Vis., vol. 23, no. 2, pp. 199–211, 2005.
- [7] K. Du, J. Lu, H. Sekiya, Y. Sun, and T. Yahagi, "Post-processing for restoring edges and removing artifacts of low bit rates wavelet-based image," in *Proc. Int. Symp. Intell. Signal Process. Commun.*, Dec. 2006, pp. 943–946.
- [8] T. Kartalov, Z. A. Ivanovski, L. Panovski, and L. J. Karam, "An adaptive pocs algorithm for compression artifacts removal," in *Proc. 9th Int. Symp. Signal Process. Appl.*, 2007, pp. 1–4.
- [9] T. Tillo and G. Olmo, "Data-dependent pre- and postprocessing multiple description coding of images," *IEEE Trans. Image Process.*, vol. 16, no. 5, pp. 1269–1280, May 2007.
- [10] P. Weiss, L. Blanc-Feraud, T. Andre, and M. Antonini, "Compression artifacts reduction using variational methods: Algorithms and experimental study," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Mar./Apr. 2008, pp. 1173–1176.
- [11] K. Du, H. Han, and G. Wang, "A new algorithm for removing compression artifacts of wavelet-based image," in *Proc. IEEE Int. Conf. Comput. Sci. Autom. Eng.*, vol. 1. Jun. 2011, pp. 336–340.
- [12] C. Jung, L. Jiao, H. Qi, and T. Sun, "Image deblocking via sparse representation," *Signal Process., Image Commun.*, vol. 27, no. 6, pp. 663–677, 2012.
- [13] A. Zakhor, "Iterative procedures for reduction of blocking effects in transform image coding," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 2, no. 1, pp. 91–95, Mar. 1992.
- [14] Y. Yang, N. P. Galatsanos, and A. K. Katsaggelos, "Regularized reconstruction to reduce blocking artifacts of block discrete cosine transform compressed images," *IEEE Trans. Circuits Syst. for Video Technol.*, vol. 3, no. 6, pp. 421–432, Dec. 1993.
 [15] A. Nosratinia, "Enhancement of JPEG-compressed images by
- [15] A. Nosratinia, "Enhancement of JPEG-compressed images by re-application of JPEG," J. VLSI Signal Process. Syst. Signal, Image Video Technol., vol. 27, nos. 1–2, pp. 69–79, 2001.
- [16] A. Nosratinia, "Postprocessing of JPEG-2000 images to remove compression artifacts," *IEEE Signal Process. Lett.*, vol. 10, no. 10, pp. 296–299, Oct. 2003.
- [17] A. Foi, V. Katkovnik, and K. Egiazarian, "Pointwise shape-adaptive DCT for high-quality denoising and deblocking of grayscale and color images," *IEEE Trans. Image Process.*, vol. 16, no. 5, pp. 1395–1411, May 2007.

- [18] X. Zhang, R. Xiong, X. Fan, S. Ma, and W. Gao, "Compression artifact reduction by overlapped-block transform coefficient estimation with block similarity," *IEEE Trans. Image Process.*, vol. 22, no. 12, pp. 4613–4626, Dec. 2013.
- [19] G. Zhai, W. Lin, J. Cai, X. Yang, and W. Zhang, "Efficient quadtree based block-shift filtering for deblocking and deringing," *J. Vis. Commun. Image Represent.*, vol. 20, no. 8, pp. 595–607, Nov. 2009.
- [20] Y. Kwon, K. I. Kim, J. Tompkin, J. H. Kim, and C. Theobalt, "Efficient learning of image super-resolution and compression artifact removal with semi-local Gaussian processes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 9, pp. 1792–1805, Sep. 2015.
- [21] M.-Y. Shen and C.-C. J. Kuo, "Review of postprocessing techniques for compression artifact removal," J. Vis. Commun. Image Represent., vol. 9, no. 1, pp. 2–14, Mar. 1998.
- [22] M. V. Afonso, J.-M. Bioucas-Dias, and M. A. T. Figueiredo, "Fast image recovery using variable splitting and constrained optimization," *IEEE Trans. Image Process.*, vol. 19, no. 9, pp. 2345–2356, Sep. 2010.
- [23] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jan. 2011.
- [24] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-andplay priors for model based reconstruction," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Dec. 2013, pp. 945–948.
- [25] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.
- [26] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, Dec. 2006.
- [27] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [28] G. K. Wallace, "The JPEG still picture compression standard," *IEEE Trans. Consum. Electron.*, vol. 38, no. 1, pp. xviii–xxxiv, Feb. 1992.
- [29] C. Christopoulos, A. Skodras, and T. Ebrahimi, "The JPEG2000 still image coding system: An overview," *IEEE Trans. Consum. Electron.*, vol. 46, no. 4, pp. 1103–1127, Nov. 2000.
- [30] G. J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand, "Overview of the high efficiency video coding (HEVC) standard," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 22, no. 12, pp. 1649–1668, Dec. 2012.
- [31] J. Lainema, F. Bossen, W.-J. Han, J. Min, and K. Ugur, "Intra coding of the HEVC standard," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 22, no. 12, pp. 1792–1801, Dec. 2012.
- [32] T. Nguyen and D. Marpe, "Performance analysis of HEVC-based intra coding for still image compression," in *Proc. Picture Coding Symp.*, May 2012, pp. 233–236.
- [33] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D, Nonlinear Phenomena*, vol. 60, nos. 1–4, pp. 259–268, 1992.
- [34] J. Schmidhuber, "Deep learning in neural networks: An overview," *Neural Netw.*, vol. 61, pp. 85–117, Jan. 2015.
- [35] A. Gersho and R. M. Gray, Vector Quantization and Signal Compression. Boston, MA, USA: Kluwer, 1992.
- [36] L. Schuchman, "Dither signals and their effect on quantization noise," *IEEE Trans. Commun. Technol.*, vol. 12, no. 4, pp. 162–165, Dec. 1964.
- [37] D. Sun and W.-K. Cham, "Postprocessing of low bit-rate block DCT coded images based on a fields of experts prior," *IEEE Trans. Image Process.*, vol. 16, no. 11, pp. 2743–2751, Nov. 2007.
- [38] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [39] Kakadu Software 7.7, Accessed on Feb. 2016. [Online]. Available: http://www.kakadusoftware.com
- [40] F. Bellard. BPG 0.9.6, Accessed on Oct. 2015. [Online]. Available: http://bellard.org/bpg/



Yehuda Dar received the B.Sc. degree in computer engineering and the M.Sc. degree in electrical engineering from the Technion–Israel Institute of Technology, in 2008 and 2014, respectively, where he is currently pursuing the Ph.D. degree with the Department of Computer Science. His research interests are in signal and image/video processing, in particular, theory and applications for image and video compression, and their relation to inverse problems and sparse representations.



Michael Elad (F'12) received the B.Sc., M.Sc., and D.Sc. from the Department of Electrical Engineering, Technion–Israel Institute of Technology (Technion), Israel, in 1986, 1988, and 1997, respectively. Since 2003, he has been a Faculty Member with the Computer-Science Department, Technion, where he has also been a Full Professor since 2010. He works in the field of signal and image processing, specializing in inverse problems and sparse representations. He received numerous teaching awards, and was a recipient of the 2008 and 2015 Henri Taub

Prizes for academic excellence, and the 2010 Hershel-Rich prize for innovation. He has been serving as the Editor-in-Chief of the *SIAM Journal on Imaging Sciences* since 2016.



Alfred M. Bruckstein received the B.Sc. and M.Sc. degrees from the Technion–Israel Institute of Technology (Technion), Haifa, in 1976 and 1980, respectively and the Ph.D. degree in electrical engineering from Stanford University, CA, in 1984. Since 1984, he has been with Technion, where he holds of the Ollendorff Chair in Science. His current research interests are in swarm/ant robotics, image and signal processing, analysis and synthesis, pattern recognition, and various aspects of applied geometry. He authored and co-authored over 150 journal papers

in the fields of interest mentioned. He is an SIAM Fellow for contributions to signal processing, image analysis, and ant robotics. He received the 2014 SIAG/Imaging Science Prize (for the From Sparse Solutions of Systems of Equations to Sparse Modeling of Signals and Images with David Donoho and Michael Elad).



Raja Giryes received the B.Sc., M.Sc. (supervision by Prof. M. Elad and Prof. Y. C. Eldar), and Ph.D. (supervision by Prof. M. Elad) degrees from the Department of Computer Science, Technion–Israel Institute of Technology (Technion), Haifa, in 2007, 2009, and 2014, respectively. He held a post-doctoral position with the Computer Science Department, Technion (from 2013 to 2014) and the Laboratory of Prof. G. Sapiro, Duke University, Durham, USA (2014 to 2015). He is currently an Assistant Professor with the School of Electrical

Engineering, Tel Aviv University. His research interests lie at the intersection between signal and image processing and machine learning, and in particular, in deep learning, inverse problems, sparse representations, and signal and image modeling. He received the Maof Prize for excellent young faculty (2016 - 2019), the VATAT Scholarship for excellent postdoctoral fellowship (2014 - 2015), the Intel Research and Excellence Award (2005, 2013), the Excellence in Signal Processing Award (ESPA) from Texas Instruments (2008), and was part of the Azrieli Fellows program (2010-2013).