

On the optimization of the satellite imaging chain

Mikael Carlavan, Laure Blanc-Féraud, Marc Antonini, Carole Thiebaut,
Christophe Latry and Yves Bobichon

Abstract

In this paper, we focus on the global optimization of the satellite imaging chain. The theoretical analysis of the satellite imaging chain optimization is a difficult problem that needs a lot of approximations. In order to consider the complex real satellite imaging chain, we propose to address this problem numerically and we present, based on numerical experiments, techniques to optimize the quality of the reconstructed final image. We first focus on the common question of the position of the restoration step in the imaging chain, that is on-board before coding or on-ground after coding. Then, we present several methods to remove the coding artifacts inherent in wavelet based coder schemes. From these numerical results we propose a new satellite imaging chain and we show visual and rate-distortion results on a real satellite image.

Index Terms

Global optimization, satellite imaging chain, coding, image restoration

I. INTRODUCTION

The composition of a classical satellite imaging chain is represented Fig. 1. This representation is simplified for more clarity and does not take into account any satellite specific preprocessing

M. Carlavan and L. Blanc-Féraud are with MORPHEME team, joint project between INRIA, CNRS and the University of Nice-Sophia Antipolis; M. Antonini is with I3S laboratory UMR7271 University of Nice Sophia-Antipolis and CNRS, CS 40121 - 06903 Sophia Antipolis, France (e-mail: carlavan@i3s.unice.fr, blancf@i3s.unice.fr, am@i3s.unice.fr).

C. Thiebaut and C. Latry are with CNES, 18 avenue Edouard Belin, 31401 Toulouse, France (e-mail: Carole.Thiebaut@cnes.fr, Christophe.Latry@cnes.fr).

Y. Bobichon is with Thales Alenia Space, 100 boulevard du Midi, 06156 Cannes la Bocca, France (e-mail: Yves.Bobichon@thalesalieniaspace.com).

operations performed on board. As shown on Fig. 1, a satellite imaging chain is mainly composed by three parts: The acquisition which captures and samples the scene, the compression which reduces the volume of data of the image before storing it in the on-board memory (compression is also required for a fast transmission) and the restoration which attenuates the degradations collected during the coding and the acquisition chains. For the rest of this paper, we assume that the transmission does not alter the quality of the image.

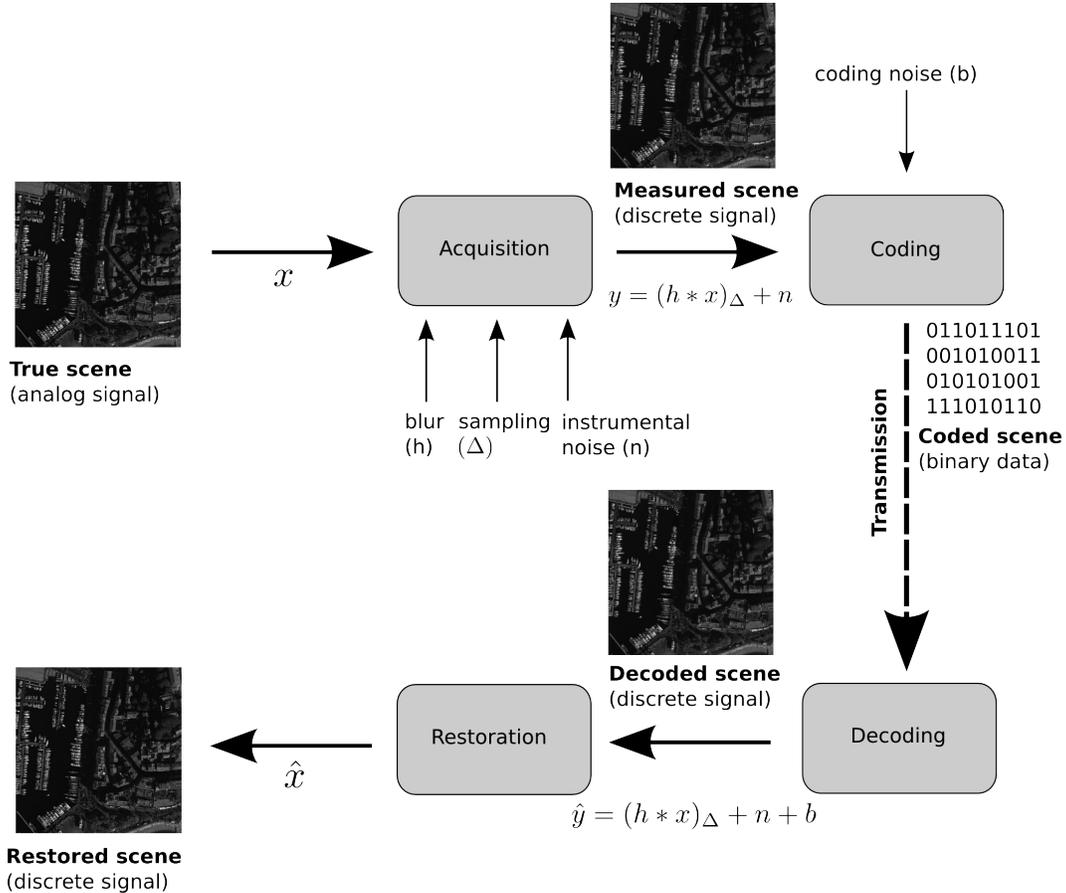


Fig. 1. Processing chain for satellite imaging.

Usually, the acquisition is fixed and fully characterized by the optics of the satellite and the electronics of the chain. The parameters of this part are therefore fixed such that the acquisition cannot be tuned so much to improve the quality of the final image. On the contrary, the compression and the restoration have many more degrees of freedom. The choice of the parameters of these two parts is therefore of crucial importance in the control of the satellite

imaging chain. This is the problem of global optimization.

Formally, the global optimization of a satellite imaging chain is the problem of designing the optimal coding and restoration algorithms which give the best final image knowing the parameters of the chain (sampling, blur, noise, coding rate, ...). This problem is very challenging and has been little investigated so far. Some works have been done in designing an optimal coding/decoding structure [1] which takes into account the characteristics of the imaging chain, or an efficient restoration method which deals with the quantizing noise [2]. But to the best of our knowledge, the study of the global system optimization, which includes both compression and restoration, has not devoted much work.

To formulate this specific problem, we consider the imaging chain showed Fig. 1. Let x be the reference or target image which is the closest discrete representation of the true analog scene that we can obtain (this image can be seen as the sampling of the convolution of the analog image with a target point spread function), y the instrumental one and \hat{x} the restored image at the output of the chain. The coding/decoding and restoration are respectively denoted by the operators C and T . We denote by $D(x, \hat{x})$ some measure of the distance between the reference image and the restored one and by $R(C(y))$ some measure of the coding rate of the coded image. The restored image can be expressed as a function of the coding and the restoration by $\hat{x} = T(C(y))$. The problem of global optimization consists in finding the optimal C^* and T^* such that

$$\begin{aligned} C^*, T^* = & \arg \min_{C, T} E [D(x, T(C(y)))] , \\ & \text{subject to } C, T \\ & R(C(y)) \leq R_c \end{aligned} \quad (1)$$

where R_c is the target coding rate and E is the expectation operator with respect to the distribution law of x . Solving problem (1) is very difficult in many aspects. Firstly, problem (1) searches for the optimal coder and restoration among all available techniques, which is not tractable. Second, even if the coding and restoration methods are given and perfectly known, an analytic expression of the global distortion is usually not available and depends on the distance measure D and on the knowledge of the real unknown image x (or its statistics). As we can see, the problem (1) is difficult to solve in a general context. The main contribution of this paper is to show that this problem can however be simplified in the specific case of D being the well-known mean square error (MSE).

From the results of [3], we show here that, in the case of the MSE criterion, the global joint-optimization problem (1) can be replaced by two disjoint optimization problems. More precisely, the original image x should first be optimally estimated from the instrumental image y and, second, this estimation should be optimally coded. This result is particularly interesting in the case of satellite imaging as it addresses the position of the estimation process in the chain, that is the position of the restoration with respect to the coding step. We will present in this paper numerical experiments on this aspect.

One difficulty of the optimization of the imaging chain is also the presence of coding artifacts in the reconstructed image. We also address this issue in the paper and we present techniques to treat this problem of coding noise removal. Finally, the last contribution of the paper is to present a new satellite imaging chain. The particularity of this chain is that it deduces from the conclusions obtained on each problematic (i.e. the position of the restoration and how to deal with the coding noise) and is therefore optimized on these aspects.

The paper is organized as follows. We present, in Section II and III, numerical experiments to improve the quality of the final image. The purpose of these parts is to bring leads, mainly based on numerical results, to open problems in the design of a satellite imaging chain such as the position of the restoration in the chain and how to deal with the coding noise. From the obtained results we propose a new imaging chain in Section IV and we present reconstruction results on a real satellite image. Finally, Section V concludes the paper and presents perspectives for future works.

The data presented in this paper are provided by the French Space Agence (CNES¹) and are simulations of the post PLEIADES-HR satellite. We will then focus only on the imaging chain system of the CNES but the methods we propose are more general and can be easily extended to the characteristics of other systems.

II. THE POSITION OF THE RESTORATION IN THE IMAGING CHAIN

The actual restoration method used by the CNES only deals with the blur and the additive Gaussian noise of the instrument [4]. It actually does not take into account the fact that the decoded image is also deteriorated with coding noise. This restoration is therefore also suitable

¹Centre National d'Etudes Spatiales

to be used on-board just before coding on the instrumental image, as this image perfectly matches the image formation model considered by the restoration.

In this part, we are thus considering the question of the position of the restoration in the satellite imaging chain, i.e. on-board² before coding (see Fig. 2) or as usual on-ground after coding (see Fig. 1).

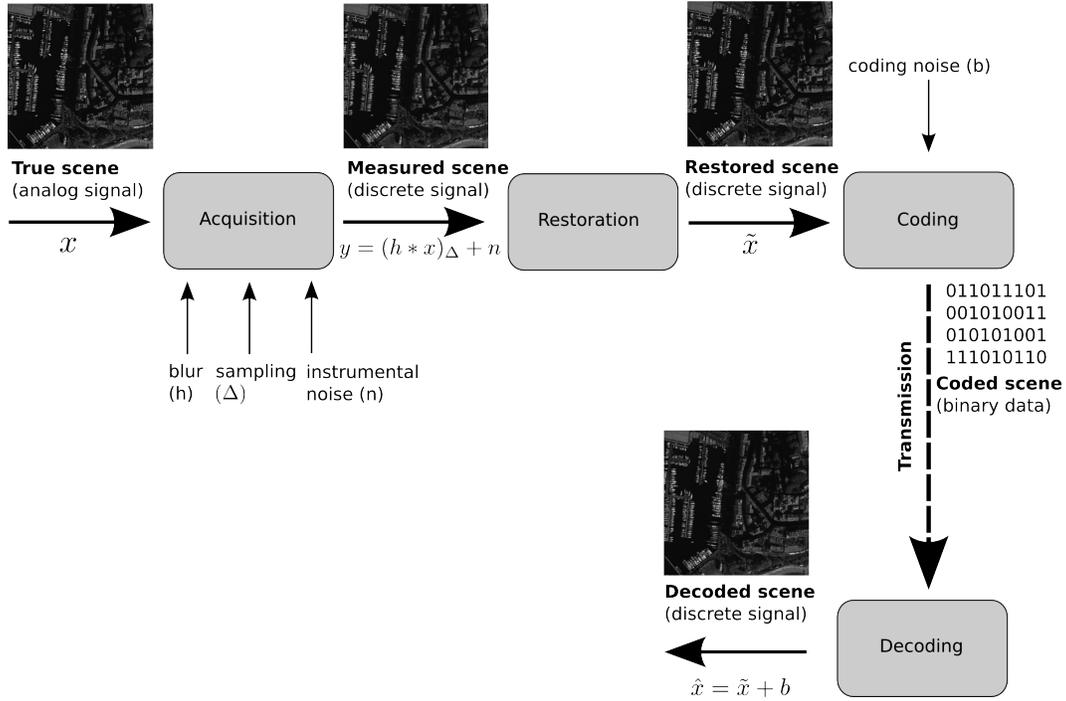


Fig. 2. On-board restoration based satellite imaging chain.

A. Introducing the on-board restoration

As mentioned in the introduction, the initial global optimization problem consists in finding the optimal coding C^* and restoration T^* such that

$$\begin{aligned}
 C^*, T^* = & \arg \min E [D(x, T(C(y)))] . \\
 \text{subject to } & C, T \\
 & R(C(y)) \leq R_c
 \end{aligned} \tag{2}$$

²The purpose of this part is to give a methodological answer to the position of the restoration in the chain; so we assume that we have at our disposal enough on-board computing resources such that an on-board restoration can be implemented

A similar problem has been formalized in [3] for the study of optimal noisy source coding. The main result of [3] states that the global distortion, if measured by the mean square error (MSE), can be separated in two terms. More precisely, the authors of [3] showed that

$$E \left[\|x - T(C(y))\|_2^2 \right] = E \left[\|x - E[x|y]\|_2^2 \right] + E \left[\|E[x|y] - T(C(y))\|_2^2 \right], \quad (3)$$

where $E[x|y]$ is the conditional expectation of the original image x knowing the noisy one y . The image $E[x|y]$ can be seen as the best estimator (which minimizes the MSE) of the original image x from y . This image does not depend on the restoration or the compression technique used. So the minimal distortion D^* then writes [3]

$$D^* = E \left[\|x - E[x|y]\|_2^2 \right] + \min_{C, T} E \left[\|E[x|y] - T(C(y))\|_2^2 \right]. \quad (4)$$

subject to C, T

From (4), we propose to use the MSE as the global measure criterion and we rewrite the optimization problem (2) as

$$C^*, T^* = \arg \min_{C, T} E \left[\|E[x|y] - T(C(y))\|_2^2 \right]. \quad (5)$$

subject to C, T
 $R(C(y)) \leq R_c$

We see that the global distortion can be expressed and optimized with respect to the image $E[x|y]$ instead of the original image x . Note that the problem (5) is not simpler to solve as the computation of the image $E[x|y]$ is usually not accessible. This formulation may however be interesting to address the common question of the position of the restoration in the design of the satellite imaging chain (that is before or after coding).

As mentioned previously, the image $E[x|y]$ represents the restoration of the true image x from the instrumental one y . It is then very tempting to think that this ideal image is actually the result of the restoration T , moved on-board of satellite (i.e. before coding), and that we can replace $E[x|y]$ in (5) by $T(y)$ such that the global optimization problem can also be written

$$C^*, T^* = \arg \min_{C, T} E \left[\|T(y) - C(T(y))\|_2^2 \right]. \quad (6)$$

subject to C, T
 $R(C(T(y))) \leq R_c$

It is certain that the problem (6) is slightly different from the initial optimization problem (2) but problem (6) is easier to treat as each variable can almost be optimized separately. If

T is fixed, then problem (6) looks for the optimal coder C^* which minimizes the coding error under the constraint that the coding rate does not exceed the target coding rate. This problem is well-known and referred as the coding rate-allocation problem [6] which has been addressed a lot in the coding community [7], [8], [9] and references therein.

To be clear, the global joint optimization problem (2) is very difficult to address. But, in our opinion, we believe that moving the restoration on-board could take benefit of the characteristics of the actual restoration and coding algorithms. So one way (but again this is not the only one) to address the problem of global joint optimization (2) is to use an on-board restoration such that the global optimization problem can be splitted in two independent ones. The first problem is to optimize the on-board restoration such that it is close to $E[x|y]$. The second problem is to design a coder C which minimizes the coding error. As mentioned previously, the latter has been the focus of intense work in the imaging community. So the difficulty here is to evaluate how close to $E[x|y]$ is $T(y)$. As the ideal image $E[x|y]$ depends on the original image x and is therefore not accessible, we will simulate several state-of-the-art restoration algorithms and observe their impact on the global distortion and on the quality of the reconstructed image. This is the focus of the next part.

B. Comparison of on-board and on-ground chains

We are considering the on-board chain showed Fig. 2 in comparison to the on-ground one illustrated Fig. 1. For the simulation, the coding step is fixed and is based on the method proposed in [10]. This method is very similar to the one implemented in the PLEIADES-HR satellite [4]. For each chain (on-board and on-ground), the restoration will be performed in two steps as follows. First, a direct deconvolution will be performed using the target point spread function (PSF) provided by the CNES. The second step is the denoising step. A wavelet packet decomposition [11] is usually employed to fit the frequential characteristics of the deconvolved noise [4]. However, another important point to take into account for an efficient denoising is the decrease rate of reconstruction error from the M largest wavelet coefficients [12]. The faster the reconstruction error decreases, the better the denoising is. And on this point, a wavelet packet transform is not optimal [13].

We propose here to perform the denoising using a variant of the wavelet transform named the Shearlet transform [14]. A wavelet transform can be represented using a matrix with dyadic shifts

and dilations as coefficients. It is classically extended to the two dimensional case using separable wavelets which process each dimension of the image independently. The matrix representation of a two dimensional wavelet transform is therefore diagonal. The Shearlet transform presented in [14] proposes instead to use a non-diagonal matrix and more specifically considers a “shear” matrix. A shear matrix is a matrix that combines operations along its rows and columns. This implies that a Shearlet transform uses combinations of shifts and dilations of each dimension of the image. This offers the ability to capture oriented details and is, among the contourlets [15] and the curvelets [16], an optimal transform (in term of reconstruction error decreasing rate with respect to the number of retained coefficients) for the representation of images [12]. A deconvolution method based on the Shearlet transform has been proposed in [12]. We will therefore compare the method [12] to the current state-of-the-art restoration methods such as the ForWarRD method [17], which performs a deconvolution followed by a regularization in both the Fourier and wavelet domains, or the method based on a Stein block thresholding [18] which performs the regularization in the Vaguelet-Wavelet domain followed by an adaptive block thresholding. Note that to be coherent with the current restoration technique used by the CNES, we only focus here on the methods which decompose the restoration in a direct deconvolution followed by a threshold operation of some sparse representation. We did not include the methods based on a variational framework such as [19].

We simulate these two chains on the image presented Fig. 3 and we compare these results to the reconstructed images provided by the CNES. The acquisition parameters have been provided by the CNES and simulate a SNR of 30-100 (each value defines the power of the noise at a certain target luminance). The restoration method used by the CNES is similar to the method proposed in [11] and consists in a direct deconvolution followed by a thresholding in a wavelet packet basis.

From numerical experiments, the threshold parameters which minimize the MSE of each method then been computed using an exhaustive. In this simulation, the original image x is known and the MSE can thus be computed. Note that in a real environment, unbiased estimator of the MSE exists and does not not require the knowledge of the true image [20].

We evaluate the reconstruction results both visually and using the peak signal-to-noise ratio



Fig. 3. Reference image of Cannes harbour (12 bits, resolution of 30 cm, 1024×1024 pixels).

(PSNR) criterion defined for 12 bits dynamic images as

$$PSNR(x, \hat{x}) = 20 \log_{10} \left(\frac{4095}{\frac{1}{N} \|x - \hat{x}\|_2} \right), \quad (7)$$

where N is the number of pixels. We show visual results only for the case of 2.5 bits/pixel as at high coding rate (4.0 bits/pixel) the reconstructed images of the on-board and on-ground chains appear perceptually identical and have an almost equal PSNR. This can be explained by the fact that at this rate, the compression is almost lossless and the coding step could then be “omitted”, making both chains equivalent.

The comparison of these two chains in a rate-distortion sense is given Fig. 4. Visual results are given Fig. 5 to 6. We can see that for every restoration techniques, an on-board chain always performs better in term of PSNR. At low coding rate, the difference between the two chains reaches almost 1 dB . We can check on Fig. 5 for example that the on-board chain gives edges which are slightly more blurred than the on-ground chain (particularly visible around the edges of buildings). This is due to the fact that the edges of the image have been enhanced by the deconvolution. The high frequency subbands require then more bits to be properly encoded.

It is actually difficult to conclude on the difference between the two chains as they both give

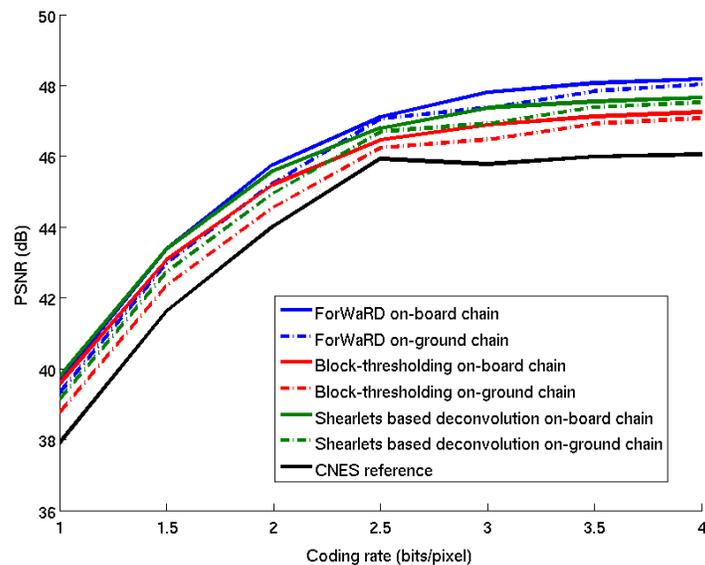


Fig. 4. Rate-distortion of on-board and on-ground chains for different restoration techniques compared to the method currently used by the CNES.

similar results, although the on-ground one seems to perform better on low intensity areas. For example, on Fig. 6, we see that the on-board chain reconstructs an image which is more blurred (see the small square element at the bottom of the figure) than the one we would have obtained with an on-ground chain (see also Fig. 5). The on-board chain presents however the advantage to separate the process of coding noise removal and we will exploit this ability later in Section IV.

Visually, the Stein block thresholding restoration technique [18] does not give satisfying results and tends to oversmooth the image. If we observe the reconstructed images, we can verify that all the small details are lost. The ForWaRD method [17] seems also to suffer from the same behavior and provides slightly smooth reconstructed images. The method based on the Shearlets [12] seems to be slightly superior in term of image quality. This method give satisfying results and recover the small details of the image without giving too many artifacts. A deeper evaluation of the reconstructed images, by image analysis experts, may be however required to confirm this result.

Finally, we see that many coding artifacts still appear in the reconstructed images. This

phenomenon is particularly visible on the reconstruction results of the on-board chain as the coding noise is not treated at all by this chain. The on-board chain may be therefore penalized by the presence of these artifacts, so we present in the next part some of the state-of-the-art processing methods to reduce these coding artifacts.

We can also observe that each restoration technique outperforms the restoration technique used by the CNES in terms of PSNR. For a coding rate of 2.5 bits/pixel, the improvement, in terms of PSNR, of these methods over the method of the CNES varies between 1 and 1.5 *dB*. Note that the PSNR of the method used by the CNES is almost constant after the coding rate of 2.5 bits/pixel as this technique leaves some residual noise to give the image a physical sense. This residual noise simulates the instrumental noise that one obtains at the output of a sensor. This phenomenon only appears from 2.5 bits/pixel, as at this rate the encoder starts to efficiently encode the instrumental noise instead of removing it. Also note that this image characteristic is highly appreciated by image analysis experts. We will use this characteristic to propose a technique to remove the coding artifacts inherent in wavelet-based compression systems. This is the focus of the next part.

III. QUANTIZING ARTIFACTS REMOVAL

In this part, we briefly describe the state-of-the-art of quantization noise removal methods. Note that this part is only a brief review of quantization noise removal techniques and we will discuss the integration of these techniques in the satellite chain in the last part of the paper.

A. Variational methods for denoising quantization noise

The quantization noise is very specific and visually comes out as structured artifacts which represent the responses of the wavelets. This phenomenon appears as oscillations, basically near the edges of the image (ringing artifacts), when an initial non-null wavelet coefficient has been set to zero by the quantization. Several methods based on a variational framework have been recently proposed [21], [22], [2] to tackle the problem of quantization noise removal. Techniques [22] and [2] are actually very similar and, consequently, we only present the methods proposed in [21] and [22]. These methods rely on the minimization of an *a priori* named the total variation (TV) [23].

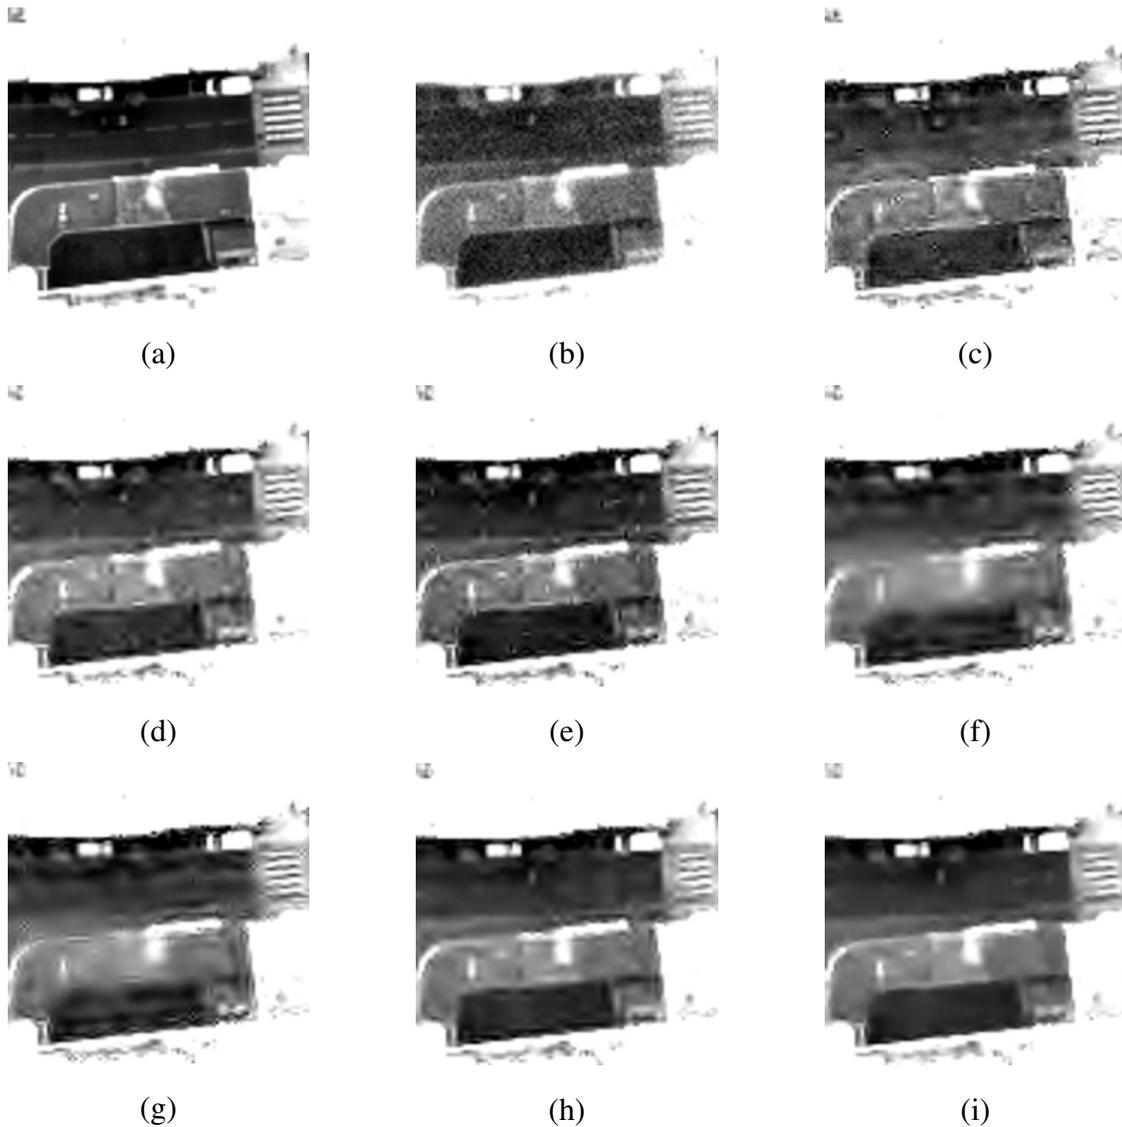


Fig. 5. Visual comparison of on-board and on-ground chains. Displayed images have a size of 200×200 pixels. (a) is the reference image, (b) is the instrumental image (output of the acquisition, $PSNR = 32.69$ dB), (c) is the reconstructed image provided by the CNES ($PSNR = 45.93$ dB), (d) and (e) are the reconstructed images respectively from the Shearlets based on-board ($PSNR = 46.80$ dB) and on-ground ($PSNR = 46.69$ dB) chains, (f) and (g) are the reconstructed images respectively from the block thresholding based on-board ($PSNR = 46.46$ dB) and on-ground ($PSNR = 46.24$ dB) chains, (h) and (i) are the reconstructed images respectively from the ForWarRD based on-board ($PSNR = 47.11$ dB) and on-ground ($PSNR = 47.05$ dB) chains. The target rate is 2.5 bits/pixel. The image range has been extended to point up the image reconstruction artifacts.

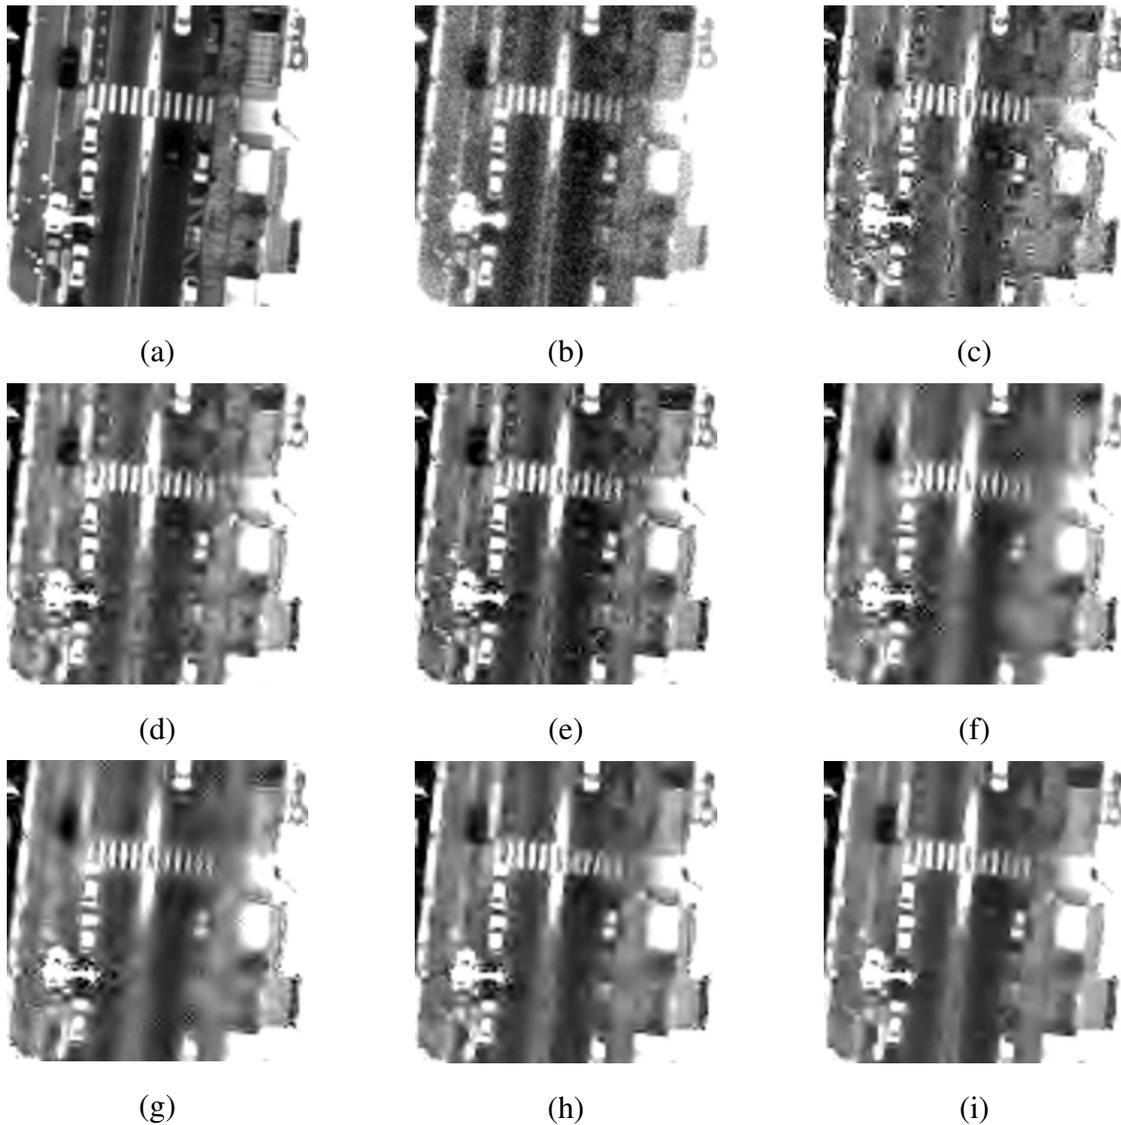


Fig. 6. Visual comparison of on-board and on-ground chains. Displayed images have a size of 200×200 pixels. (a) is the reference image, (b) is the instrumental image (output of the acquisition, $PSNR = 32.69$ dB), (c) is the reconstructed image provided by the CNES ($PSNR = 45.93$ dB), (d) and (e) are the reconstructed images respectively from the Shearlets based on-board ($PSNR = 46.80$ dB) and on-ground ($PSNR = 46.69$ dB) chains, (f) and (g) are the reconstructed images respectively from the block thresholding based on-board ($PSNR = 46.46$ dB) and on-ground ($PSNR = 46.24$ dB) chains, (h) and (i) are the reconstructed images respectively from the ForWarRD based on-board ($PSNR = 47.11$ dB) and on-ground ($PSNR = 47.05$ dB) chains. The target rate is 2.5 bits/pixel. The image range has been extended to point up the image reconstruction artifacts.

The TV prior supposes that an image can be modeled as a smooth function with discontinuities across curves. The oscillations created by the artifacts cannot therefore be considered to be natural and do not belong to an image. The simple way to remove these variations is to minimize the l^1 -norm of the gradient of the image, namely the TV, such that these oscillations are replaced by smooth homogeneous regions. Both methods [21] and [22] could globally be formalized as the following minimization problem

$$\begin{aligned} \hat{x} = \quad & \arg \min \quad \|\nabla x\|_1 \quad , \\ & \text{subject to} \quad x \in K \end{aligned} \quad (8)$$

where \hat{x} is the denoised image and K is a set that constrains the reconstructed image. The authors of [22] proposed to define the set K such that it constrains the error between the observed and the reconstructed wavelet coefficients to be bounded by the boundaries of the true quantization error.

Let x_0 be the image to encode, W the wavelet transform of the coder and \tilde{x} the decoded image. We denote \mathcal{Q} the set of all possible output quantized values $\mathcal{Q} = \{q_k; k \in \mathbb{Z}, q_0 = 0\}$ and b_k, b_{k+1} ($b_{k+1} > b_k$) the boundaries of each quantization interval such that

$$(W\tilde{x})_i = q_k, \quad \text{if } b_k \leq (Wx_0)_i < b_{k+1}, \quad \forall i \in \{0, \dots, N-1\}. \quad (9)$$

We have

$$b_k - q_k \leq (Wx_0)_i - (W\tilde{x})_i < b_{k+1} - q_k, \quad \forall i \in \{0, \dots, N-1\}. \quad (10)$$

We set $\alpha_i = b_k - q_k$ and $\beta_i = b_{k+1} - q_k$; the authors of [22] proposed to define K as the following hypercube

$$K = \left\{ x \in \mathbb{R}^N, \alpha_i \leq (Wx)_i - (W\tilde{x})_i < \beta_i, \forall i \in \{0, \dots, N-1\} \right\}. \quad (11)$$

The true quantization error is unknown as the original image x_0 is not accessible. The bounds α_i and β_i of this error can however be estimated from the reconstructed image and the knowledge of the quantizing model. So the originality of the method proposed in [22] consists in minimizing the TV of the reconstructed image such that the quantization error belongs to the intervals defined by the boundaries (10) of the true quantization error. The method proposed in [21] is slightly different and constrains the wavelet coefficients that have not been set to zero by the quantization to remain identical. They define the set K as

$$K = \left\{ x \in \mathbb{R}^N, (Wx)_i = (W\tilde{x})_i, \forall i \in M \right\}, \quad (12)$$

where M is the set of coefficients coordinates that have not been set to zero by the quantizing

$$M = \{i \in \{0, \dots, N - 1\}, |(W\tilde{x})_i| > 0\}. \quad (13)$$

The idea of the method proposed in [21] is to reconstruct the small coefficients that have been set to zero by the quantization. The method relies on the fact that the minimization of the TV creates flat regions which need small wavelet coefficients to be represented. The presence of the constraint (12) is to ensure that only the small coefficients are updated and that the large quantized coefficients, which are likely to be close to the original ones, remain unchanged.

A comparison of the two presented methods is given at the end of this part. We will see however that the flat homogeneous regions created by these methods are not natural in the sense that they cannot be interpreted physically. The problem of quantization noise removal is actually very difficult to address. The main difficulty lies in the fact that the quantization noise is highly correlated to the signal source and cannot be modeled using classical probability distributions. We present in the next part methods to improve the statistical properties of the quantization noise.

B. Dithering methods for removing quantization artifacts

A dithering technique consists in inserting a noise prior to quantizing to improve the statistics of the quantization error. Among the dithering techniques, we focus on the subtractive dithering technique proposed in [24] whose particularity is to subtract the added noise after quantizing. Let w be the wavelet coefficients of the image to quantize and Q a uniform scalar quantizer of quantizing step Δ . The quantized coefficients \tilde{w} are obtained by applying the quantizing operator Q on the wavelet coefficients w noised by the dithering noise v

$$\tilde{w} = Q(w + v). \quad (14)$$

As mentioned previously, a subtractive dithering scheme also subtracts the dithering noise after the quantizing. The final wavelet coefficients z are

$$z = \tilde{w} - v = Q(w + v) - v. \quad (15)$$

The authors of [24] showed that the global error $\epsilon = z - w$ of a subtractive dithering system is independent of the system source and is distributed uniformly if the dither noise v can be

expressed as a summation of rectangular probability density functions [25]. More precisely, the probability density function p_ϵ of the global error ϵ can be expressed as [24]

$$p_\epsilon(\epsilon) = \Delta \Pi_\Delta(\epsilon) \cdot [W_\Delta * p_w * p_v](-\epsilon), \quad (16)$$

where p_v is the probability density function of the dithering noise v , p_w is the probability density function of the wavelet coefficients w to quantize and

$$\Pi_\Delta(\epsilon) = \begin{cases} \frac{1}{\Delta} & \text{if } -\frac{\Delta}{2} < \epsilon \leq \frac{\Delta}{2}, \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

$$W_\Delta(\epsilon) = \sum_{k=-\infty}^{+\infty} \delta(\epsilon - k\Delta). \quad (18)$$

The characteristic function (defined as the Fourier transform of the probability density function) is usually easier for studying the correlation properties. It writes [26]

$$\begin{aligned} P_\epsilon(u) &= \text{sinc}(u) * [W_{\frac{1}{\Delta}}(-u)P_w(-u)P_v(-u)] \\ &= \sum_{k=-\infty}^{+\infty} \text{sinc}\left(u - \frac{k}{\Delta}\right) P_w\left(-\frac{k}{\Delta}\right) P_v\left(-\frac{k}{\Delta}\right) \\ &= \text{sinc}(u) + \sum_{k=-\infty, k \neq 0}^{+\infty} \text{sinc}\left(u - \frac{k}{\Delta}\right) P_w\left(-\frac{k}{\Delta}\right) P_v\left(-\frac{k}{\Delta}\right). \end{aligned} \quad (19)$$

where

$$\text{sinc}(u) = \begin{cases} \frac{\sin(\pi\Delta u)}{\pi\Delta u}, & \text{if } u \neq 0 \\ 1, & \text{otherwise} \end{cases}. \quad (20)$$

From (19), we see that the global error ϵ is independent of the source w only if the characteristic function P_ϵ is reduced to $\text{sinc}(u)$. This is verified if the characteristic function of the dither noise P_v cancels for each $\frac{k}{\Delta}$, $\forall k \in \mathbb{Z}^*$ [24]

$$P_v\left(\frac{k}{\Delta}\right) = 0, \quad \forall k \in \mathbb{Z}^*. \quad (21)$$

As mentioned by [25], Eq. (21) is satisfied if the dithering noise can be expressed as a summation of rectangular probability density functions [25]. In that case, the global error ϵ reduces to an independent and uniformly distributed noise.

This is an encouraging result as it implies that an on-board restoration coupled with a subtractive dithering scheme will result in a restored image with a residual noise which is

independent of the original image. This residual noise can then be interpreted physically as the instrumental noise of the sensor. This aspect of residual noise is very important as it is one of the features sought by the CNES for the design of restoration methods [27]. We will discuss this aspect later as this is the basis of the proposed imaging chain described in Section IV.

Finally, we would like to mention also the dithering technique proposed by [28]. The method proposed in [28] is slightly different from the classical dithering techniques as it is more focus on the reconstruction of the original wavelet subbands rather than improving the statistics of the quantizing noise. More precisely, the main result of [28] is that the probability density function of a wavelet subband can be recovered exactly (assuming we know the parameters of its model) from its quantized version by adding a dithering noise v to the quantized coefficients. As previously, we denote by w an original (i.e. prior to quantizing) wavelet subband and \tilde{w} the corresponding quantized subband. The authors of [28] proposed to model a wavelet subband w (each subband can be treated separately) by a Laplace distribution [29]

$$p_w(w) = \frac{\lambda}{2} e^{-\lambda|w|}, \quad (22)$$

where λ is the scale parameter that can be estimated using classical estimation techniques such that least-squares minimization methods or maximum-likelihood estimations. Similarly to (9), we have

$$\tilde{w} = q_k, \quad \text{if } b_k \leq w < b_{k+1}. \quad (23)$$

Using the wavelet subband model (22), we can express the probability density function of a quantized wavelet subband

$$p_{\tilde{w}}(\tilde{w} = q_k) = \begin{cases} \frac{1}{2} (e^{-\lambda b_k} - e^{-\lambda b_{k+1}}), & \text{if } k \geq 1 \\ 1 - \frac{1}{2} (e^{-\lambda b_0} - e^{-\lambda b_1}), & \text{if } k = 0 \\ \frac{1}{2} (e^{\lambda b_{k+1}} - e^{\lambda b_k}), & \text{if } k \leq -1. \end{cases} \quad (24)$$

The output wavelet subband z is given by

$$z = \tilde{w} + v, \quad (25)$$

where \tilde{w} is the quantized wavelet subband and v the dithering noise. The output wavelet subband probability density function can be expressed using the law of total probability [28]

$$p_z(z) = \sum_k p_{z|\tilde{w}}(z|\tilde{w} = q_k) p_{\tilde{w}}(\tilde{w} = q_k), \quad (26)$$

where

$$p_{z|\tilde{w}}(z|\tilde{w} = q_k) = p_{v|\tilde{w}}(v = z - q_k|\tilde{w} = q_k), \quad (27)$$

is the probability density function of the dither noise v knowing the quantized values \tilde{w} . The authors of [28] showed that the choices

$$p_{v|\tilde{w}}(v|\tilde{w} = q_k, k \neq 0) = \begin{cases} \frac{1}{\alpha_k} e^{-\text{sign}(q_k)\hat{\lambda}v}, & \text{if } (b_k - q_k) \leq v < (b_{k+1} - q_k) \\ 0, & \text{otherwise} \end{cases} \quad (28)$$

$$p_{v|\tilde{w}}(v|\tilde{w} = 0) = \begin{cases} \frac{1}{\alpha_0} e^{-\hat{\lambda}|v|}, & \text{if } b_0 > v > b_1 \\ 0, & \text{otherwise} \end{cases}, \quad (29)$$

with α_k being normalization constants and $\hat{\lambda}$ an estimated value of the scale parameter λ , lead to the original wavelet subband probability density function, under the condition that $\hat{\lambda} = \lambda$ [28]

$$\begin{aligned} p_z(z) &= \sum_k p_{z|\tilde{w}}(z|\tilde{w} = q_k) p_y(\tilde{w} = q_k) \\ &= \sum_{k \leq -1} \frac{1}{\alpha_k} e^{\lambda(z-q_k)} \frac{1}{2} (e^{\lambda b_{k+1}} - e^{\lambda b_k}) \mathbb{1}(b_k \leq z < q_{k+1}) \\ &\quad + \frac{1}{\alpha_0} e^{-\lambda|z|} \left(1 - \frac{1}{2} (e^{\lambda b_0} - e^{-\lambda b_1}) \right) \mathbb{1}(b_0 \leq z < b_1) \\ &\quad + \sum_{k \geq 1} \frac{1}{\alpha_k} e^{-\lambda(z-q_k)} \frac{1}{2} (e^{-\lambda b_k} - e^{-\lambda b_{k+1}}) \mathbb{1}(b_k \leq z < q_{k+1}) \\ &= \frac{\lambda}{2} e^{-\lambda|z|} = p_w(z) \end{aligned} \quad (30)$$

Even if the reconstructed and original subbands will numerically differ, this technique will remove the undesirable observed artifacts, due to the quantization, by filling in the blanks. The fact that we also add dither noise on the null coefficients may also provide the residual noise appreciated by the image analysis experts.

C. Comparison of removal methods for quantization artifacts

We simulate the behavior of the presented quantization removal methods directly on a coded version of the reference (i.e. without any blur or noise) satellite image shown Fig. 3. The simulation of the complete imaging chain including these techniques is done in the next part. To perform a fair comparison, the image will be coded using the biorthogonal 9/7 wavelet

transform [30] followed by the quantizer described in [26]. As a consequence, the method [28] has been adapted to this choice. For the subtractive dithering method [25], we simulated a uniform dithering noise to limit the power of the residual noise. We only provide visual results as common criteria such as PSNR do not take into account the appreciated physical perception of residual noise.

The results are given Fig. 7 to 8. Visually, we immediately see that the techniques based on the minimization of the TV create large smooth homogenous regions and remove the small details of the image. This effect is known as the *cartoon* effect. These flat regions are not considered to be natural for a remote sensing image and are really not appreciated by image analysis experts who clearly prefer a deterioration that can be interpreted physically. As explained previously, this is for example the case of an unstructured residual noise. The subtractive dithering system and the method proposed in [28] give good visual results. Both images are well reconstructed and do not present common artifacts such as ringing or blurry edges. The quality of the image reconstructed with the subtractive dithering system actually seems slightly better, particularly on the small details of the image (cars and zebras). As expected, these methods leave a residual noise on the reconstructed image which simulates the instrumental noise that image expert analysis are used to observe.

IV. PROPOSED IMAGING CHAIN

In the previous section, we showed that the dithering techniques may be very interesting to remove the structured artifacts of the coding step. As we have also mentioned in Section III, these techniques leave a uniform residual noise which is highly appreciated from the image analysis experts as it can be interpreted physically. More precisely, an ideal final image (as defined by image analysis experts) should own a residual blur characterized by a target PSF [5] along with a uniform residual noise with a fixed standard deviation [27].

We also presented in Section II an on-board restoration technique which gives an image with a residual noise (whose power is very small in comparison to the power of the residual noise obtained from the dithering technique) and a residual blur fully characterized by the target PSF. If we combine these two techniques, i.e. if we use an on-board restoration coupled with a subtractive dithering technique, the image obtained at the output of the chain will then present a residual noise (coming from the dithering technique) with the blur of the target PSF (coming from the

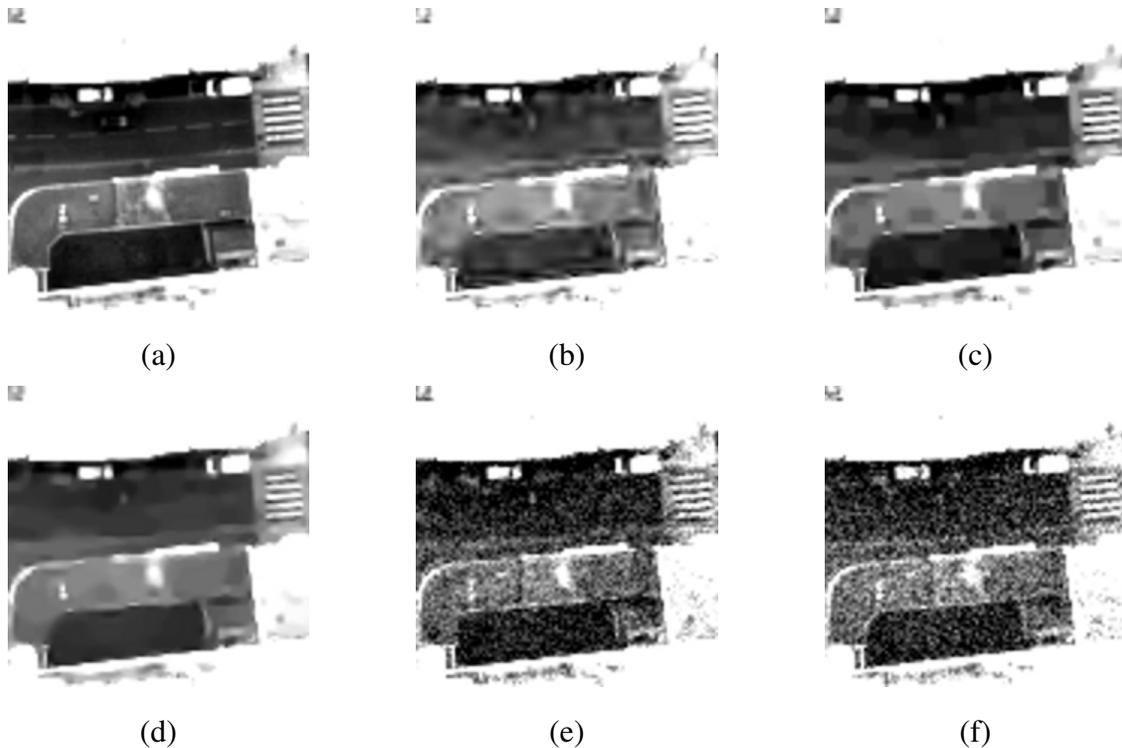


Fig. 7. Visual comparison of quantizing removal techniques. Displayed images have a size of 200×200 pixels. (a) is the reference image, (b) is the decompressed image, (c) is the image obtained using the post-processing technique proposed in [21], and (d) is the image obtained using the post-processing technique proposed in [22], (e) is the image reconstructed using the post-processing dithering technique proposed in [28], (f) is the image reconstructed using the subtractive dithering technique [26] with an uniform dithering noise. The target rate is 2.5 bits/pixel. The image range has been extended to point up the image reconstruction artifacts.

on-board restoration). And as mentioned previously, a final image with such characteristics is the objective of image analysis experts as it can be interpreted as the direct output of an ideal instrument.

From this remark, we propose the imaging chain shown Fig. 9.

This chain includes the on-board restoration [12] based on the Shearlets transform and the subtractive dithering technique [25] to decorrelate the quantizing noise. Note that, in this chain, the quantizer follows the model described in [26] to respect the subtractive dithering scheme hypothesis. The coding step is then decomposed in a 3-levels CDF 9/7 wavelet transform followed by an explicit quantization of the wavelet coefficients and an entropy encoding of the quantized coefficients. The results of the proposed imaging chain are given Fig. 10 to 13. Like previously,

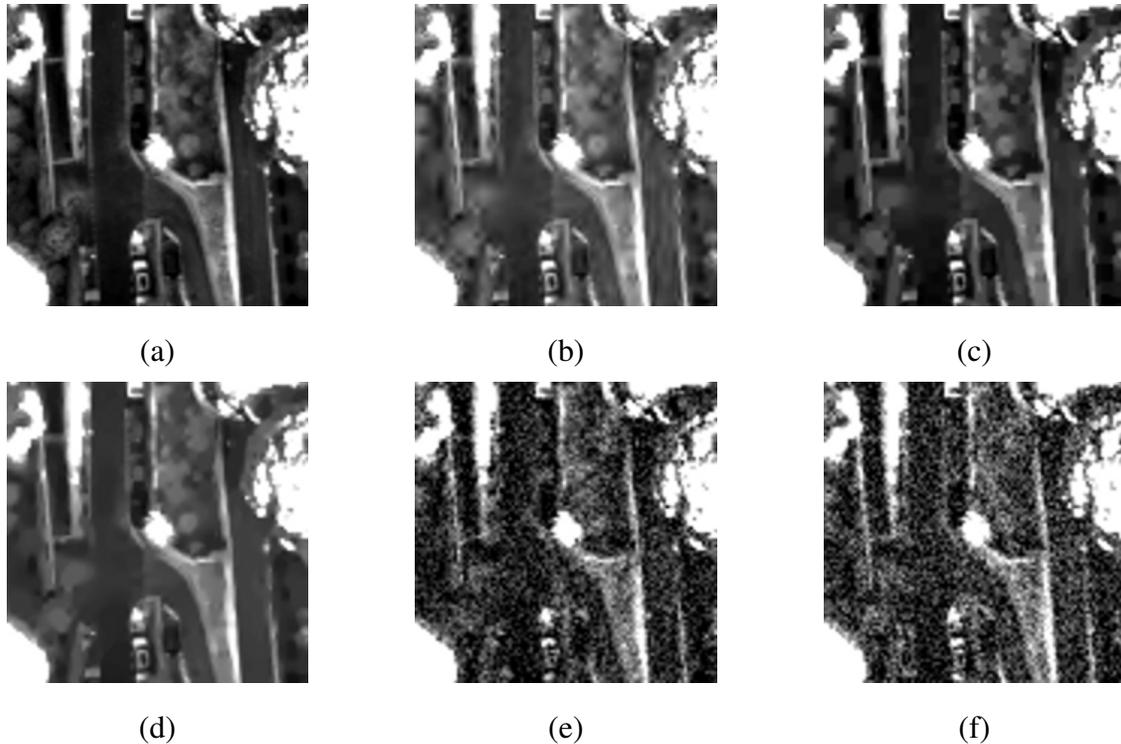


Fig. 8. Visual comparison of quantizing removal techniques. Displayed images have a size of 200×200 pixels. (a) is the reference image, (b) is the decompressed image, (c) is the image obtained using the post-processing technique proposed in [21], and (d) is the image obtained using the post-processing technique proposed in [22], (e) is the image reconstructed using the post-processing dithering technique proposed in [28], (f) is the image reconstructed using the subtractive dithering technique [26] with an uniform dithering noise. The target rate is 2.5 bits/pixel. The image range has been extended to point up the image reconstruction artifacts.

the simulated SNR of this simulation is 30-100.

We immediately see that the reconstructed images with the proposed chain do not present any common wavelet compression artifacts (see Fig. 10 and 11), that we observed on the reconstructed image provided by the CNES. They exhibit instead an unstructured residual noise which is visually similar to the noise obtained on the instrumental image at the output of the acquisition chain. This is particularly visible on the dark zones of the reconstructed image, see Fig. 11. It is clear that the proposed chain tends to replace one type of residual noise (wavelet compression artifacts) by another one. The obtained residual noise is however better appreciated by image analysis experts as it can be interpreted physically. More precisely, the proposed imaging chain produces a reconstructed image which owns the two characteristics of an ideal image: Blur with

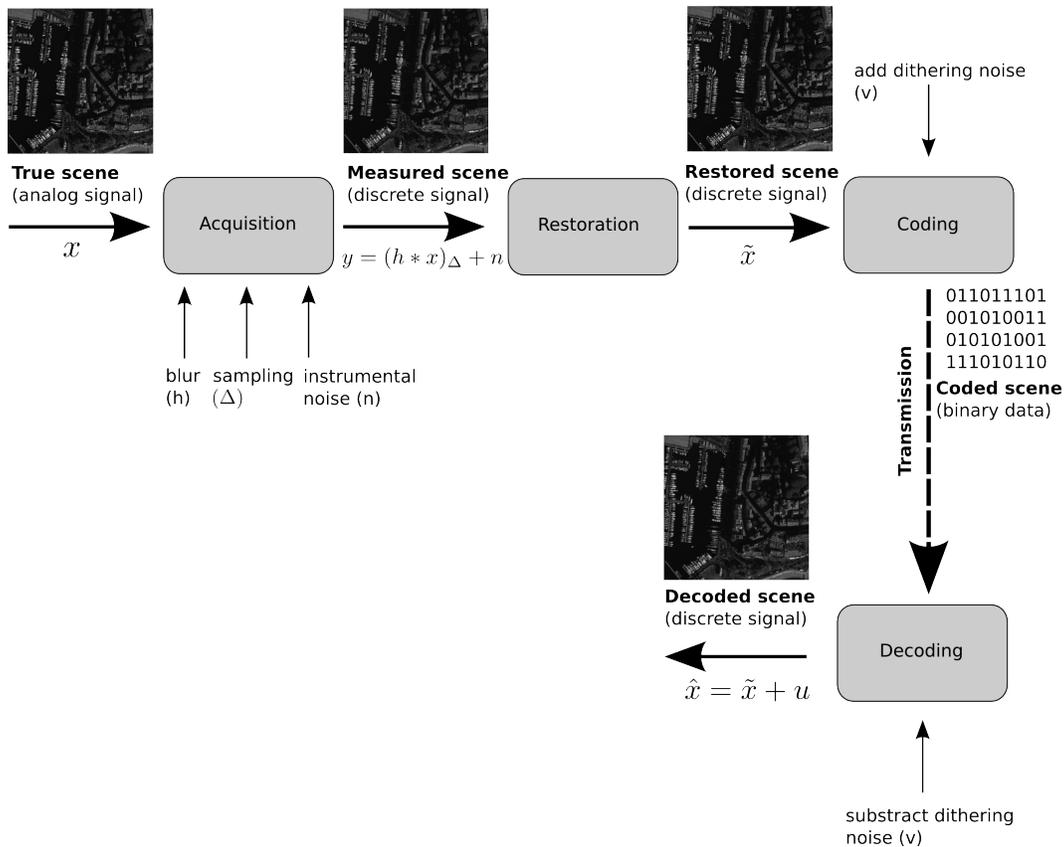


Fig. 9. Proposed satellite imaging chain

the target PSF (obtained by the on-board restoration) and a residual unstructured noise [27]. The drawback of the proposed method is that the standard deviation of the residual noise is function of the quantizing step while it should be constant for all coding rates. Consequently, for a low coding rate, the proposed chain gives an image which is more noisy than the instrumental one. It gives however very interesting results for high coding rates as shown by Fig. 12 to 13. Further works need thus to be done on this aspect.

V. CONCLUSIONS

In this paper, we considered the satellite imaging chain optimization problem. We presented numerical results which showed that the quality of the reconstructed image can be improved if one concedes several changes on the usual design of imaging chain. The first one would be to move the restoration step on-board of satellite, prior to compression. This allows to reconstruct

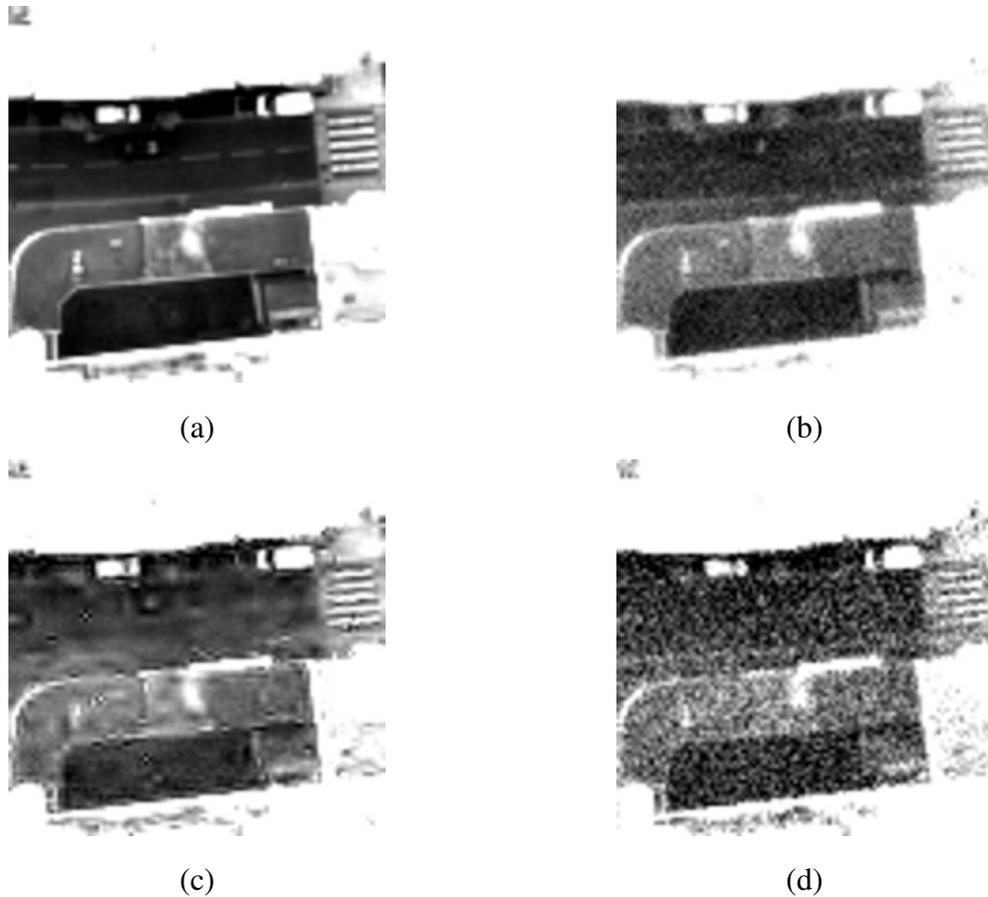


Fig. 10. Visual comparison of the proposed and the current imaging chains. Displayed images have a size of 200×200 pixels. (a) is the reference image, (b) is the instrumental image, (c) is the decompressed and restored image provided by the CNES, (d) is the reconstructed image from the Shearlets based on-board chain followed by a subtractive dithering scheme. The target rate is 2.5 bits/pixel. The image range has been extended to point up the image reconstruction artifacts.

an image with less reconstruction artifacts, specially on shadows zones. The second deals with the problem of coding noise removal. From the results we presented, we concluded that the current state-of-the-art algorithms do not give competitive results and that the best option may be to use dithering techniques to transform this structured coding noise into an unstructured uniform residual noise. This residual noise is very interesting as it simulates the noise obtained at the output of the instrument, and this property is highly appreciated from photo interpreters. From these conclusions, we proposed a new imaging chain based on an on-board restoration coupled with a subtractive dithering technique. We showed results on a real satellite data. We compared the results of the proposed chain with the ones obtained with the current satellite imaging chain.

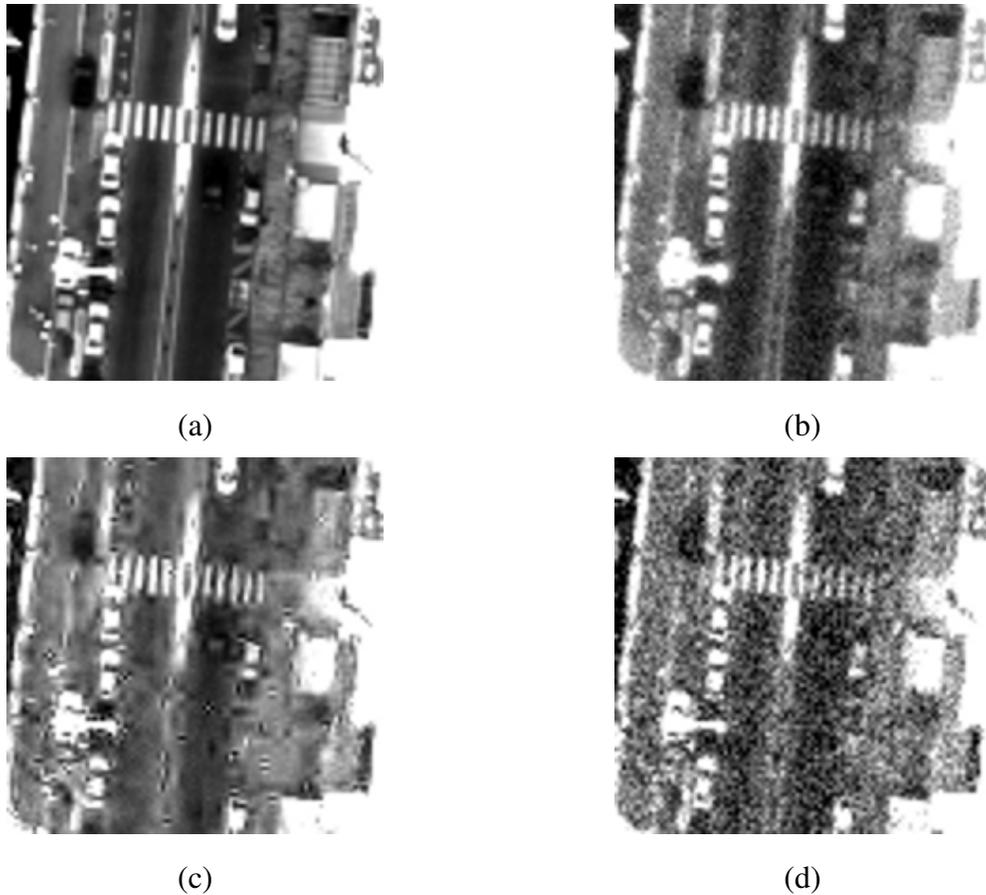


Fig. 11. Visual comparison of the proposed and the current imaging chains. Displayed images have a size of 200×200 pixels. (a) is the reference image, (b) is the instrumental image, (c) is the decompressed and restored image provided by the CNES, (d) is the reconstructed image from the Shearlets based on-board chain followed by a subtractive dithering scheme. The target rate is 2.5 bits/pixel. The image range has been extended to point up the image reconstruction artifacts.

We showed that the proposed chain gives interesting results and may be particularly efficient at medium and high coding rates (around 3.0 bits/pixel and more). The particularity of the proposed imaging chain is that the final image is fully characterized by the target blur (specified by the CNES) and a residual unstructured noise. Such feature is interesting for images analysis experts since classical defects of the compression and restoration steps do not appear in the final image, such that these two steps appear then almost transparent in the chain. A drawback of the proposed method is that the power of this residual noise depends on the target coding rate. At low coding rate (like 2.5 bits/pixel), the final image appears to be more noisy than the instrumental image and is therefore difficult to exploit. Future works will be focussed to extend our approach such

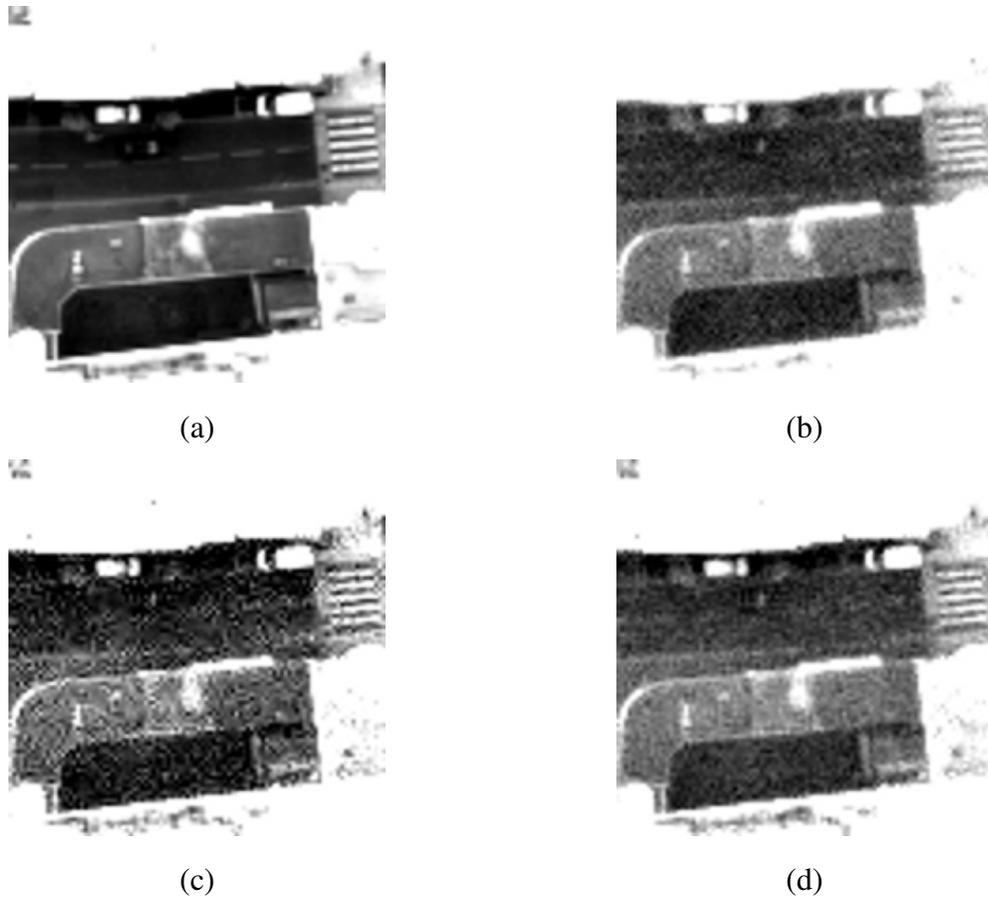


Fig. 12. Visual comparison of the proposed and the current imaging chains. Displayed images have a size of 200×200 pixels. (a) is the reference image, (b) is the instrumental image, (c) is the decompressed and restored image provided by the CNES, (d) is the reconstructed image from the Shearlets based on-board chain followed by a subtractive dithering scheme. The target rate is 4.0 bits/pixel. The image range has been extended to point up the image reconstruction artifacts.

that we get competitive results at low coding rates.

REFERENCES

- [1] C. Parisot, M. Antonini, M. Barlaud, S. Tramini, C. Latry and C. Lambert-Nebout. Optimization of the Joint Coding/decoding Structure. *IEEE International Conference on Image Processing*, 2001.
- [2] S. Tramini, M. Antonini, M. Barlaud and G. Aubert. Optimal joint decoding/deblurring method for optical images. *IEEE International Conference on Image Processing*, 1:381–385, 1999.
- [3] J. Wolf, J. Ziv. Transmission of noisy information to a noisy receiver with minimum distortion. *IEEE Transactions on Information Theory*, 16(4):406–411, 1970.
- [4] P. Lier, C. Valorge and X. Briottet. Imagerie spatiale : des principes d’acquisition au traitement des images optiques pour l’observation de la Terre. *Editions Cépaduès*, 2008.

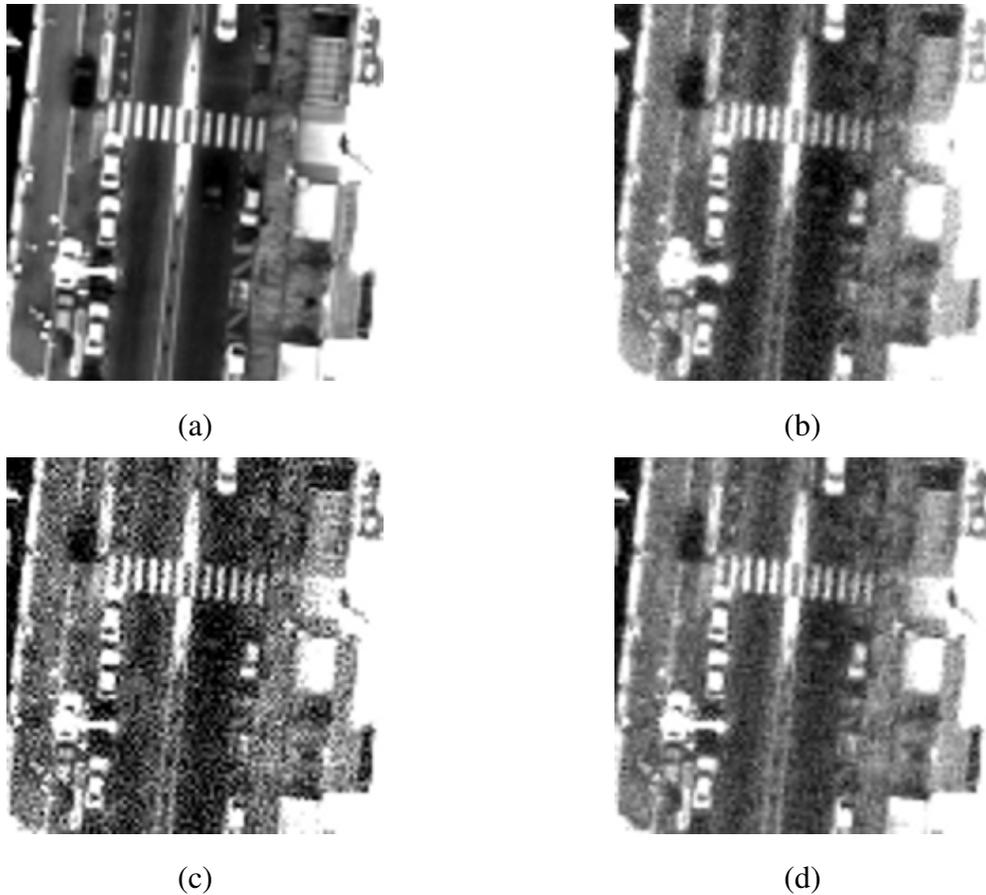


Fig. 13. Visual comparison of the proposed and the current imaging chains. Displayed images have a size of 200×200 pixels. (a) is the reference image, (b) is the instrumental image, (c) is the decompressed and restored image provided by the CNES, (d) is the reconstructed image from the Shearlets based on-board chain followed by a subtractive dithering scheme. The target rate is 4.0 bits/pixel. The image range has been extended to point up the image reconstruction artifacts.

- [5] C. Lambert-Nebout, C. Latry, G.A. Moury, C. Parisot, M. Antonini, M. Barlaud. On-board optical image compression for future high-resolution remote sensing systems. *Proc. SPIE Applications of Digital Image Processing*, 2000.
- [6] C.E. Shannon. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
- [7] M. Antonini, M. Barlaud, P. Mathieu, P. and I. Daubechies. Image coding using wavelet transform. *IEEE Transactions on Image Processing*, 1(2):205–220, 1992.
- [8] T. Berger. Rate Distortion Theory: A Mathematical Basis for Data Compression. *Prentice-Hall*, 1971.
- [9] A. Ortega, K. Ramchandran. Rate-distortion methods for image and video compression. *IEEE Signal Processing Magazine*, 15(6):23–50, 1998.
- [10] Consultative Committee for Space Data Systems. Image Data Compression. *Blue Book*, 2005.
- [11] J. Kalifa, S. Mallat, B. Rouge. Deconvolution by thresholding in mirror wavelet bases. *IEEE Transactions on Image Processing*, 12(4):446–457, 2003.

- [12] V.M. Patel, G.R. Easley, D.M. Healy. Shearlet-Based Deconvolution. *IEEE Transactions on Image Processing*, 18(12):2673–2685, 2009.
- [13] S. Mallat. A wavelet tour of signal processing. *Academic Press*, 2008.
- [14] D. Labate, W.-Q. Lim, G. Kutyniok and G. Weiss. Sparse multidimensional representation using shearlets. *Proc. SPIE, Wavelet Applications in Signal and Image Processing XI*, 5914:254–262, 2005.
- [15] M.N. Do and M. Vetterli. The contourlet transform: an efficient directional multiresolution image representation. *IEEE Transactions on Image Processing*, 14(12):2091–2106, 2005.
- [16] E.J. Candès, L. Demanet, D.L. Donoho, L. Ying. Fast discrete curvelet transforms. *Multiscale Model. Simul.*, 5(3):861–899, 2006.
- [17] R. Neelamani, H. Choi, R. Baraniuk. ForWaRD: Fourier-Wavelet Regularized Deconvolution for Ill-Conditioned Systems. *IEEE Transactions on Signal Processing*, 52(2):418–433, 2004.
- [18] C. Chesneau, J. Fadili, J.-L. Starck. Stein block thresholding for wavelet-based image deconvolution. *Electronic Journal of Statistics*, 4:415–435, 2010.
- [19] J. Bect, L. Blanc-Féraud, G. Aubert, A. Chambolle. A l1-unified variational framework for image restoration. *European Conference on Computer Vision*, 2004.
- [20] S. Ramani and T. Blu and M. Unser. Monte-Carlo SURE: A Black-Box Optimization of Regularization Parameters for General Denoising Algorithms. *IEEE Transactions on Image Processing*, 17(9):1540–1554, Sep. 2008.
- [21] S. Durand and J. Froment. Reconstruction Of Wavelet Coefficients Using Total Variation Minimization. *SIAM Journal on Scientific Computing*, 24:1754–1767, 2003.
- [22] P. Weiss, L. Blanc-Feraud, T. Andre, and M. Antonini. Compression artifacts reduction using variational methods: algorithms and experimental study. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2008.
- [23] L. Rudin, S. Osher, E. Fatemi. Nonlinear Total Variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
- [24] L. Schuchman. Dither Signals and Their Effect on Quantization Noise. *IEEE Transactions on Communication Technology*, 12(4):162–165, 1964.
- [25] R.A. Wannamaker, S.P. Lipshitz, J. Vanderkooy, J.N. Wright. A theory of nonsubtractive dither. *IEEE Transactions on Signal Processing*, 48(2):499–516, 2000.
- [26] S.P. Lipshitz, R.A. Wannamaker and J. Vanderkooy. Quantization and Dither: A Theoretical Survey. *Journal of the Audio Engineering Society*, 40(5):355–375, 1992.
- [27] P. Dherete, and B. Rouge. Image de-blurring and application to SPOT5 THR satellite imaging. *IEEE International Geoscience and Remote Sensing Symposium*, 1:318–320, 2003.
- [28] M.C. Stamm, K.J.R. Liu. Anti-Forensics of Digital Image Compression. *IEEE Transactions on Information Forensics and Security*, 6(3):1050–1065, 2011.
- [29] J. Li, R.M. Gray. Text and picture segmentation by the distribution analysis of wavelet coefficients. *IEEE International Conference on Image Processing*, 3:790–794, 1998.
- [30] A. Cohen, I. Daubechies, J.-C. Feauveau. Biorthogonal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics*, 45(5):485–560, 1992.