

Dequantizing image orientation

IEEE Transactions on Image Processing, Vol. 11, No. 10, pp. 1129-1140, October 2002.

Agnès Desolneux, Saïd Ladjal, Lionel Moisan, Jean-Michel Morel

CMLA, ENS Cachan,

61 avenue du président Wilson,

94235 Cachan cedex - France

phone : 33 1 47 40 59 45 / fax : 33 1 47 40 59 01

e-mail : {desolneu,ladjal,moisan,morel}@cmla.ens-cachan.fr

Abstract

We address the problem of computing a local orientation map in a digital image. We show that standard image grey level quantization causes a strong bias in the repartition of orientations, hindering any accurate geometric analysis of the image. In continuation, a simple dequantization algorithm is proposed, which maintains all of the image information and transforms the quantization noise in a nearby Gaussian white noise (we actually prove that only Gaussian noise can maintain isotropy of orientations). Mathematical arguments are used to show that this results in the restoration of a high quality image isotropy. In contrast with other classical methods, it turns out that this property can be obtained without smoothing the image or increasing the SNR. As an application, it is shown in the experimental section that, thanks to this dequantization of orientations, such geometric algorithms as the detection of nonlocal alignments can be performed efficiently. We also point out similar improvements of orientation quality when our dequantization method is applied to aliased images.

Keywords

Dequantization, orientation map, alignment detection.

I. INTRODUCTION

Let $u(x)$ be a grey level image, where x denotes the pixel and $u(x)$ is a real value. Most natural (non synthetic) images are generated in the following way : a source image s is assumed to be of infinite resolution. A band limited optical smoothing is performed on s , yielding a smoothed version $k * s$. By Shannon-Whittaker theory, the band-limited image can be sampled on a regular and fine enough grid. Let us denote by Π the Dirac Comb of this grid. Then u is roughly obtained as $u = (s * k) \cdot \Pi$, which yields the discrete, digital image. According to Shannon-Whittaker Theorem, $s * k$ can be recovered from u by the so called Shannon interpolation, using a basis of sinc functions. Actually, this model is significantly idealized, since other operations result in a substantial image degradation, namely a white photonic and/or electronic noise n , a windowing (Π is not infinite, but restricted to a rectangle) and, last but not least, a quantization Q . Thus, the realistic image model is $u = Q[(k * s) \cdot \Pi + n]$, in which we neglect the windowing effect as affecting essentially the image boundary. In this paper, we address the problem of computing accurately and in an unbiased way the orientation of the gradient of u , a number $\theta \in [0, 2\pi]$ such that $\exp(i\theta) = Du/|Du|$, where $Du = (u_x, u_y)$ denotes the image gradient. When we refer to the gradient of u , we wish to refer to the gradient of the smooth subjacent image, in as much as we consider u to be Shannon interpolable. If we assume, which is realistic

enough, that k and n are isotropic, we are led to address the effect of the quantization Q on the field of orientations. We discovered recently that this effect is strong and leads to a very biased field of orientations. It can hinder any faithful geometric analysis of the image, unless some previous restoration is performed. Before explaining how we shall address this restoration, let us give an example where this restoration is crucial in order to perform a correct geometric analysis in the image. This is a particular instance, but let it be mentioned that all probabilistic methods using local pixel interactions (e.g. Markov random field models) would suffer, knowingly or not, the same effect. We proposed recently a grouping, nonlocal, method for the detection of alignments in an image. In a few words, the principle of the method is the following ([3]). We assume that each point in the image has an orientation $\Theta(x)$ (equal to the orientation of the gradient plus $\pi/2$). We consider a segment S of aligned points in the image, with length l . Let Θ_0 denote the orientation of this segment. Assume we have observed k points on S (among the l points) having their orientation equal, according to a given precision p , to the orientation of S (i.e. such that $|\Theta(x) - \Theta_0| \leq p\pi$). If k is large enough, then we say that the segment S is meaningful (more details about this method are given in the last section). In Figure 6.b, we show all segments detected in a natural image by this method at precision $p = 1/16$. It can be visually checked that no detected segment seems to be artefactual, i.e. due to image generation. Let us now choose a precision of orientation $p = 1/64$. This precision may seem exagereate, but can be successfully used in an image with strong gradients. Figure 6.c shows the detected alignments, which are, according to our definition, highly noncasual. Clearly, such detections are artefactual and the result of image generation. After some inquiry, it turned out that the grey level quantization is responsible for such artefactual detections. Actually, this does not mean that the alignment detection is wrong, but only that the detected alignments are image generation artefacts. In Figure 6.d, we show the result of alignment detection (at precision $p = 1/64$) on the same picture, after the dequantization we propose here has been performed.

Let us therefore come back to the problem of computing a reliable orientation. The first good answer to this problem is known as the dithering method [5] which consists in adding a noise before quantization and then substracting the same noise from the quantized image. This results in decreasing the SNR of the image, but turns out to better maintain the image aspect and its isotropy under strong quantization. Unfortunately, the dithering method has been to the best of our knowledge fully abandonned in image generation devices. To summarize, image isotropy can

be restored by dithering to the cost of decreasing the SNR, but this is a degradation and should anyway be performed in the image generation process itself : this is not generally the case.

A second easy answer, much in use, consists in smoothing the image by some convolution kernel, and only retaining the orientation at points where the gradient is high and stable across scales. This is the classical “edge detection” method (see [15], [10], [2] and, for more up to date methods, [9]). There is nothing to object to this method, since at the end it retains edge points which are very local, though they are confirmed at larger scales. Now, clearly, many orientations in the image can be used to detect alignments, which are not computed at edge points : the edge points simply are a particularly good selection, but sparse. Another way, addressed recently and successfully by several authors ([12], [14]) consists in defining an orientation scale-space. Also, the affine scale space ([13], [1]) provides a way to compute a multiscale orientation of level lines. In all cases, the objective is different and wider than just computing a local orientation : the aim of these methods is to compute a multiscale orientation map which has to be considered by itself as a nonlocal analysis of the image. These methods are better than edge maps methods in the sense that they provide an orientation at all points. They are all the same not appropriate for image analysis models based on local observations (e.g. most probabilistic methods), as well as the one we outlined before. Indeed, they do not preserve the independence of points at Nyquist distance.

The solution we propose in order to dequantize the image should, according to the preceding discussion, satisfy the following requirements :

- maintain the “independence” of local observations (i.e. no smoothing),
- maintain all of the image information (thus the method must be invertible),
- give an unbiased orientation map, where quantization noise has been made isotropic.

We shall actually prove that a simple and invertible operation, namely a Shannon $(\frac{1}{2}, \frac{1}{2})$ -translation of the image, permits to remove the quantization effects on the orientation map. More precisely, we shall prove experimentally and mathematically that this translation transforms the quantization noise into a nearby Gaussian white noise. We shall also prove that all reasonable local computations of the gradient, applied to the dequantized image, yield an unbiased orientation, even at points where the gradient is small. This remains true even when the quantization step is large. As a consequence, we point out the possibility of performing the geometric analysis of an image with a very local estimate of the gradient, using therefore the full

image accuracy.

Our plan is the following. In Section 2, we consider a wide set of classical local computation methods for the gradient and show that they preserve an excellent isotropy, under the assumption that the image noise is uniform or Gaussian. We also prove a converse statement : the image orientation will be isotropic if and only if the noise is Gaussian. We analyze the bias introduced by quantization and show that its effect on orientation can be disastrous. In Section 3, we detail the proposed solution and make an accurate mathematical and practical analysis of the dequantized noise. We show that it is nearby Gaussian, therefore permitting the local computation of orientations. In Section 4, we end with some experiments.

II. Local computation of gradient and orientation

We consider a discrete grey-level image u of size $N \times M$. At each point, we can compute the gradient on a 2×2 neighbourhood (we choose the smallest possible one to preserve locality). It is defined by

$$\mathbf{G}(n, m) := \begin{pmatrix} u_x \\ u_y \end{pmatrix} := \frac{1}{2} \begin{pmatrix} X_2 + X_4 - X_1 - X_3 \\ X_1 + X_2 - X_3 - X_4 \end{pmatrix}, \quad (1)$$

where $X_1 = u(n, m)$, $X_2 = u(n + 1, m)$, $X_3 = u(n, m + 1)$ and $X_4 = u(n + 1, m + 1)$ (see Fig. 1).

X1	X2
X3	X4

Fig. 1. The four pixel values used to compute the discrete gradient.

Aside from a classical finite differences estimate of the gradient of u , (1) can be interpreted as the exact gradient $\nabla \tilde{u}(n + 1/2, m + 1/2)$, where \tilde{u} is the bilinear interpolate of u defined in $[n, n + 1] \times [m, m + 1]$ by

$$\tilde{u}(x, y) = (y - m) \left((x - n)X_4 + (1 - x + n)X_3 \right) + (1 - y + m) \left((x - n)X_2 + (1 - x + n)X_1 \right).$$

From (1), we write $R = |G|$ and define the orientation θ by

$$u_x + iu_y = R \exp(i\theta). \quad (2)$$

Note that θ is not defined when $R = 0$. Our aim will be to study the behaviour of θ as a function of the four values X_1, X_2, X_3 and X_4 . The question is to decide whether such a way of computing the orientation is valid or not (i.e. whether it gives some privilege to particular directions or not). In this section we prove that if the image u is a Gaussian white noise, then there is no bias on the orientations (this means that, at each point, all orientations have an equal probability), and that, if u is a uniform white noise, there is a small bias (orientations multiple of $\pi/4$ are slightly favoured).

A. Gaussian noise

We first show that if the image u is a Gaussian white noise, then there is no bias on the orientations.

Proposition 1: Let X_1, X_2, X_3 and X_4 be independent identically Gaussian $\mathcal{N}(0, \sigma^2)$ distributed random variables. Then θ is uniformly distributed on $[0, 2\pi]$.

Proof : The first point is to notice that if we denote $A = X_2 - X_3$ and $B = X_1 - X_4$, then A and B are independent and $u_x = \frac{A-B}{2}$ and $u_y = \frac{A+B}{2}$. Thus, from (2),

$$A + iB = R\sqrt{2} \cdot \exp \left[i(\theta - \pi/4) \right]. \quad (3)$$

Now, A is Gaussian with mean 0 and variance $2\sigma^2$ because it is the sum of two independent Gaussian random variables (with mean 0 and variance σ^2). And B is also Gaussian for the same reason. Since A and B are independent, the law of the couple (A, B) is given by the density function

$$f(a, b) = \frac{1}{4\pi\sigma^2} \exp \left(-\frac{a^2 + b^2}{4\sigma^2} \right),$$

which shows that θ is almost surely defined. Last, since f is isotropic (it only depends on the squared radius $a^2 + b^2$), we deduce that the distribution of θ is uniform on $[0, 2\pi]$. \square

Proposition 2 (Converse proposition) Let X_1, X_2, X_3 and X_4 be four independent identically distributed random variables. Assume that their common law is given by a probability density f , where f is square integrable and even. If the law of θ is uniform on $[0, 2\pi]$, then the probability density f is Gaussian.

Proof : As we did in the proof of Proposition 1, we denote $A = X_2 - X_3$ and $B = X_1 - X_4$. A and B are independent and identically distributed. They have the same density function g given by the convolution of $x \mapsto f(x)$ with itself, i.e. $g(x) = \int_{-\infty}^{+\infty} f(x-y)f(y)dy$. Since f is square

integrable, the function $x \mapsto g(x)$ is continuous. We also notice that $g(0) = \int f^2(y)dy \neq 0$ (because f is a density function). Since the law of θ is uniform on $[0, 2\pi]$, we know that the law of (A, B) , given by the density $g(x)g(y)$, only depends on $x^2 + y^2$, which can be written $g(x)g(y) = g(\sqrt{x^2 + y^2})g(0)$. In addition, g is even and never vanishes since it is a continuous function such that $g(0) \neq 0$ and $g(x)^2 = g(x\sqrt{2})g(0)$. Hence, we can consider the function \tilde{g} defined for $x \geq 0$ by $\tilde{g}(x) = \ln(g(\sqrt{x})/g(0))$. Then, we get for all $x, y \geq 0$, $\tilde{g}(x+y) = \tilde{g}(x) + \tilde{g}(y)$. Since \tilde{g} is continuous and $\tilde{g}(0) = 0$, this shows that \tilde{g} is linear. Consequently, there exists $\sigma \in \mathbb{R}$ such that for all x in \mathbb{R} , $g(x) = g(0) \exp(-x^2/2\sigma^2)$ where the constant $g(0)$ is defined by the property $\int g = 1$. Thus, the law of A (and also B) is the Gaussian distribution $\mathcal{N}(0, \sigma^2)$. We now prove that the law of the X_i is also Gaussian. Since $g = f * f$, considering the Fourier Transform, we get $\hat{g}(t) = \hat{f}(t)^2 = C \exp(-t^2\sigma^2/2)$. Thus, \hat{f} is Gaussian. Since the inverse Fourier transform of a Gaussian distribution is also Gaussian, it shows that f is the Gaussian distribution with mean 0 and variance $\sigma^2/2$. \square

This result has a strong practical consequence : if we wish to have a non biased orientation map for digital image, we **must** process the image in such a way that its noise becomes as Gaussian as possible. We shall see that it is feasible with quantization noise.

The following proposition is the generalization of Proposition 1, when the gradient is computed on a larger neighbourhood (the proof of this proposition is given in Appendix).

Proposition 3 (Generalization) Assume that the components (u_x, u_y) of the gradient are computed on n neighbouring pixels X_1, X_2, \dots, X_n , i.e. $u_x = \sum_{i=1}^n \lambda_i X_i$ and $u_y = \sum_{i=1}^n \mu_i X_i$, where λ_i and μ_i are real numbers such that $\sum_{i=1}^n \lambda_i \mu_i = 0$ and $\sum_{i=1}^n \lambda_i^2 = \sum_{i=1}^n \mu_i^2$. If the X_i are independent identically Gaussian $\mathcal{N}(0, \sigma^2)$ distributed, then the angle θ is uniformly distributed on $[0, 2\pi]$.

B. Computation of orientation on nonquantized images

In this section, we address the effect on the orientation histogram of applying the former described computation of the gradient. We shall see that the bias introduced by the method is small. It is not always realistic to assume that the local repartition of the grey levels of an image is Gaussian. Instead, we can roughly assume that the values at neighbouring points differ by a uniform random variable. Thus it is licit, or at least very indicative, to compute the orientation map of a uniform white noise, in order to have an estimate of the bias on orientation provoked by this gradient computation. Let us therefore perform the computations in the following frame-

work : consider an image whose values at pixels are independent random variables, identically uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$. We then get a small bias on the orientation θ . More precisely, we have the following proposition (the proof is given in Appendix).

Proposition 4: Let X_1, X_2, X_3 and X_4 be independent random variables, identically uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$. Then the law of θ is given by the density function g , $\pi/2$ -periodic and whose restriction to $[-\pi/4, \pi/4]$ is

$$g(\theta) = \frac{1}{12} \left((1 + \tan^2(\frac{\pi}{4} - |\theta|)) \left(2 - \tan(\frac{\pi}{4} - |\theta|) \right) \right)$$

(see Figure 2).




Fig. 2. Law of θ on $[-\pi, \pi]$ when the image is a uniform noise, and comparison with the uniform distribution on $[-\pi, \pi]$ (dotted line).

Proposition 4 shows that if the pixels of the image have independent uniformly distributed values, then the orientations are not uniformly distributed. The law of the orientation θ is given by Figure 2. It shows that the orientations multiple of $\pi/4$ are favoured. If we want to measure the bias, we can compute the relative deviation from the uniform distribution on $[-\pi, \pi]$. We get : $\text{bias} = 2\pi \times \max_{\theta} |g(\theta) - 1/2\pi| \simeq 0.047$. This shows, however, that the bias is small, about 4.7%.

C. Bias of quantization

We saw in the previous section that the way we compute the orientation at a point of the image from the gradient does not create artefacts. Now, on the contrary, we will see that the histogram of orientations in the image is very sensitive to a quantization of grey levels. Let us first consider the simplest case : a binary image. We assume that the grey level at each pixel is 0 (black) or 1

(white). Then, the orientation θ of the gradient only takes a finite number of possible values : the multiples of $\pi/4$. The binary case is an extreme case. Let us now consider the case of an image quantized on a finite number n of grey levels : $\{0, 1, 2, \dots, n-1\}$. Again, we denote $A = X_2 - X_3$ and $B = X_1 - X_4$. Then A and B have discrete values in $\{-n+1, \dots, -1, 0, 1, \dots, n-1\}$. If $A = B = 0$, θ is undefined. If $A = 0$ and $B \neq 0$, then $\theta = -\pi/4$ or $\theta = 3\pi/4$. In the other cases, we have $\tan(\theta - \pi/4) = B/A$ and consequently θ only takes a finite number of values.

Let us compute the distribution law of θ when the image is a uniform discrete noise (i.e. the X_i are independent and for all k in $\{0, \dots, n-1\}$, $P[X_i = k] = 1/n$). First, we compute the probability distribution of A (and B) :

$$\forall k \in \{-n+1, \dots, n-1\}, \quad P[A = k] = \sum_{j=0}^{n-1} P[X_3 = j] \cdot P[X_2 = k + j] = \frac{n - |k|}{n^2}.$$

Hence with probability $\frac{1}{n^2}$, $A = B = 0$ and θ is not defined. Let us now compute the probability distribution of B/A , when $A \neq 0$. For each possible discrete value $b/a \in \mathbb{Q}$ (with a and b mutually prime) of B/A , we have

$$P[B/A = b/a] = \sum_{\lambda \in \mathbb{Z}^*} P[B = \lambda b] \cdot P[A = \lambda a]. \quad (4)$$

In particular, we can compute the probability of the event $\theta = \pi/4$ (it corresponds to the event $B = 0$ and $A > 0$). Notice that, thanks to symmetries, this probability also is the probability of the events $\theta = -\pi/4$, $3\pi/4$ or $-3\pi/4$. We get

$$P[\theta = \pi/4] = \sum_{a=1}^{n-1} P[A = a] \cdot P[B = 0] = \sum_{a=1}^{n-1} \frac{n - a^3}{n} = \frac{n-1}{2n^2}.$$

Moreover, we have

$$\forall \alpha \in \mathbb{Z}, \quad \alpha \neq \frac{\pi}{4} \left[\frac{\pi}{2} \right], \quad P[\theta = \alpha] < P[\theta = \pi/4].$$

This shows that the orientations multiple of $\pi/4$ are highly favoured. On Figure 3, we plot the probability distribution of θ when the number of grey levels is $n = 6$, $n = 9$ and $n = 257$. These three cases correspond respectively to $n-1$ prime, and $n-1$ of the form 2^q . Equation (4) shows that the probability distribution of B/A is directly related to a known problem of arithmetic : how many distinct irreducible fractions of the form b/a can you make with the constraint $0 \leq b \leq a \leq N$? This problem has already been addressed, in a very similar way, in some papers of Lopez-krahe et. al (see [7] and [8]) for the detailed study of the effects induced

by lattice quantization on the histogram of the slopes of straight lines joining two points of the lattice. The main difference with our study is that we have in addition a probability distribution on the values of a and b .




Fig. 3. Probability distribution of θ for $-\pi/2 \leq \theta \leq \pi/2$, when the grey levels are respectively uniformly distributed in $\{0, 1, \dots, 5\}$ (left figure), in $\{0, 1, \dots, 8\}$ (middle figure) and in $\{0, 1, \dots, 256\}$ (right figure).

Most quantized images are not binary images, but the effect of quantization on the computation of the gradient orientation is always very significant. The reason for this is that in an image, there are usually many “flat” regions. In these regions, the grey levels take a small number of values, and consequently the orientation is very quantized.

III. ORIENTATION DEQUANTIZATION

A. The proposed solution : Fourier translation

We assume that the original signal (before quantization) is a Shannon signal (i.e. we can reconstruct the whole signal from the samples). We denote this signal by s if it is a one-dimensional signal and by u if it is a two-dimensional image :

$$s(x) = \sum_{k \in \mathbb{Z}} s(k) \cdot \text{sinc}(\pi(x - k)) \quad \text{and} \quad u(x, y) = \sum_{k, l \in \mathbb{Z}} u(k, l) \cdot \text{sinc}(\pi(x - k)) \cdot \text{sinc}(\pi(y - l)).$$

The sinc function is classically defined by $\text{sinc}(x) = \frac{\sin x}{x}$, with the convention that $\text{sinc}(0) = 1$. Now, we do not know the exact values of the $s(k)$ (resp. of the $u(k, l)$). We only have the quantized signal S (resp. U). Thus, at each point,

$$s(k) = S(k) + X_k \quad \text{or} \quad u(k, l) = U(k, l) + X_{k, l}.$$

where X_k (resp. $X_{k,l}$) is the quantization noise. In the following, we assume that the X_k (resp. $X_{k,l}$) are independent, and uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$. This independence assumption is correct above the Nyquist distance.

The proposed solution for dequantization is the following one. We replace the quantized values of the signal $S(n)$ by the Shannon interpolates $S(n + 1/2)$ and obtain

$$\begin{aligned} S(n + \frac{1}{2}) &= \sum_{k \in \mathbb{Z}} S(n+k) \frac{\sin(\pi(1/2 - k))}{\pi(1/2 - k)} \\ &= \sum_{k \in \mathbb{Z}} s(n+k) \frac{(-1)^k}{\pi(1/2 - k)} - \sum_{k \in \mathbb{Z}} X_{n+k} \frac{(-1)^k}{\pi(1/2 - k)} \\ &= s(n + \frac{1}{2}) - \sum_{k \in \mathbb{Z}} X_{n+k} \frac{(-1)^k}{\pi(1/2 - k)}. \end{aligned}$$

For the quantized image $U(n, m)$, this formula becomes

$$U(n + \frac{1}{2}, m + \frac{1}{2}) = u(n + \frac{1}{2}, m + \frac{1}{2}) - \sum_{k,l \in \mathbb{Z}} X_{n+k} X_{m+l} \frac{(-1)^k}{\pi(1/2 - k)} \cdot \frac{(-1)^l}{\pi(1/2 - l)}.$$

Remark : For a finite image of size $N \times N$, we have

$$u(x, y) = \sum_{k,l=0}^{N-1} u(k, l) \cdot \text{sinc}_d(\pi(x - k)) \cdot \text{sinc}_d(\pi(y - l)),$$

where sinc_d is the discrete version of the sinc function, defined by $\text{sinc}_d(t) = \frac{\sin(t)}{N \tan(\frac{t}{N})}$ (with the convention that $\text{sinc}_d(\pi t) = 1$ for $t = 0 [N]$).

B. Study of the dequantized noise

By the dequantization method, we aim at replacing the structured quantization noise by a noise as Gaussian as possible. We will see that, the Shannon translation being an isometry, we do not reduce or enlarge the variance of the noise. Thus, we can already claim that the method is at any rate harmless. We can of course reconstruct the original digital image by the inverse translation.

Our aim in this subsection will be to study the dequantized noise Y_1 (in 1-D) and Y_2 (in 2-D) defined by

$$Y_1 = \sum_{k \in \mathbb{Z}} \frac{(-1)^k}{\pi(1/2 - k)} X_k \quad \text{and} \quad Y_2 = \sum_{k,l \in \mathbb{Z}} \frac{(-1)^k}{\pi(1/2 - k)} \frac{(-1)^l}{\pi(1/2 - l)} X_{k,l}, \quad (5)$$

where the X_k (resp. the $X_{k,l}$) can be assumed, in a first approximation, to be independent and uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$.

Let us introduce some notations. For $k \in \mathbb{Z}$, we set

$$c_k = \frac{(-1)^k}{\pi(1/2 - k)}. \quad (6)$$

We thus have $Y_1 = \sum_{k \in \mathbb{Z}} c_k X_k$ and $Y_2 = \sum_{k,l \in \mathbb{Z}} c_k c_l X_{k,l}$. If X is a random variable uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$, then the mean and the variance of X are $E(X) = 0$ and $\text{var}(X) = 1/12$. Since $\sum c_k^2$ is convergent (it is equal to 1), the random variable series (5) defining Y_1 and Y_2 are convergent in L^2 , and we moreover have $E(Y_1) = E(Y_2) = 0$ and $\text{var}(Y_1) = \text{var}(Y_2) = 1/12$. The variance of Y_1 (and of Y_2) is the same as the variance of X_k (this can also be explained by the fact that the Fourier 1/2-translation is an isometry of L^2) and thus, we do not reduce or enlarge the variance of the noise.

On Figure 4, we show the distributions of Y_1 and Y_2 . On the same figure we plot the Gaussian distribution which has mean 0 and variance 1/12. These probability distributions seem to be very close. We shall measure this. We also notice on this figure that the distribution of Y_2 looks “more Gaussian” than the one of Y_1 . This can be qualitatively explained by the Central limit Theorem and by the fact that a larger sum of independent random variables is involved in the definition of Y_2 .




Fig. 4. Left : distribution of Y_1 and comparison with the Gaussian distribution with mean value 0 and variance 1/12 (dotted curve). Right : distribution of Y_2 and comparison with the Gaussian distribution.

B.1 Kurtosis Comparison

One way to compare the distribution of Y_1 and Y_2 to the Gaussian distribution is to compare their fourth order moment (notice that they have already same mean 0, same variance 1/12 and same third order moment 0). More precisely, we will now compare their normalized fourth order moment (called the kurtosis, see [11]).

Definition 1 (Kurtosis) The kurtosis κ of a random variable X with mean $E(X)$ and variance $\text{var}(X)$ is defined by

$$\kappa = \frac{E[(X - E(X))^4]}{\text{var}(X)^2}.$$

A classical result is that any Gaussian distribution has a kurtosis equal to 3 (it is independent of the mean and variance). We will now compute the kurtosis of the distributions of Y_1 and Y_2 . This is a very useful way to check whether a distribution is Gaussian like.

Proposition 5: Let κ_1 be the kurtosis of Y_1 and κ_2 be the kurtosis of Y_2 , then

$$\kappa_1 = \frac{13}{5} = 2.6 \quad \text{and} \quad \kappa_2 = \frac{43}{15} \simeq 2.87.$$

Proof : One way to compute the fourth order moment of Y_1 is to compute its characteristic function, denoted by Φ_{Y_1} (which is defined as the Fourier transform of the probability distribution of Y_1), and then to compute the fourth order derivative of Φ_{Y_1} at 0. We first compute

$$\Phi_{Y_1}(t) = E[e^{itY_1}] = \prod_{k \in \mathbb{Z}} E[e^{itc_k X_k}] = \prod_{k \in \mathbb{Z}} \Phi(c_k t),$$

where Φ is the characteristic function of the uniform distribution on $[-\frac{1}{2}, \frac{1}{2}]$, which is given by $\Phi(t) = \text{sinc}\left(\frac{t}{2}\right)$. Consequently, we get

$$\Phi_{Y_1}(t) = \prod_{k \in \mathbb{Z}} \text{sinc}\left(\frac{c_k t}{2}\right). \quad (7)$$

For x close to 0, we have the Taylor expansion $\text{sinc}(x) = 1 - \frac{x^2}{3!} + \frac{x^4}{5!} + O(x^6)$. Thus, for t close to 0, using also the Taylor expansion of the log function, we get

$$\log \Phi_{Y_1}(t) = -\frac{t^2}{6} S_2 + \frac{t^4}{120} S_4 - \frac{t^4}{72} S_4 + O(t^6),$$

where $S_2 = \sum_{k \in \mathbb{Z}} (c_k/2)^2$ and $S_4 = \sum_{k \in \mathbb{Z}} (c_k/2)^4$. Finally, since for x close to 0, we have the Taylor expansion $\exp(x) = 1 + x + x^2/2 + O(x^3)$, we get

$$F(t) = 1 - \frac{t^2}{6} S_2 + \frac{1}{2} \cdot \frac{t^4}{36} S_2^2 - \frac{t^4}{180} S_4 + O(t^6).$$

The fourth order moment of Y_1 is then,

$$E(Y_1^4) = \Phi_{Y_1}^{(4)}(0) = 24 \cdot \left(\frac{1}{72} S_2^2 - \frac{1}{180} S_4 \right).$$

On the other hand, we can compute S_2 and S_4 using Bernoulli numbers and the zeta function (see [4] for example), and get $S_2 = 1/4$ and $S_4 = 1/48$. Finally, we obtain

$$E(Y_1^4) = \frac{13}{720} \quad \text{and} \quad \kappa_1 = \frac{13}{720} \cdot 12^2 = \frac{13}{5}.$$

In the same way, we can compute the fourth order moment of $Y_2 = \sum_{k,l} c_k c_l X_{k,l}$,

$$E(Y_2^4) = 4! \cdot \left(\frac{1}{72} \left[\sum_{k,l \in \mathbb{Z}} \left(\frac{c_k c_l}{2} \right)^2 \right]^2 - \frac{1}{180} \sum_{k,l \in \mathbb{Z}} \left(\frac{c_k c_l}{2} \right)^4 \right).$$

$$E(Y_2^4) = 24 \cdot \left(\frac{16}{72} S_2^4 - \frac{16}{180} S_4^2 \right) = \frac{43}{2160}.$$

The variance of Y_2 is also $1/12$, and thus the kurtosis κ_2 of Y_2 is

$$\kappa_2 = \frac{43}{2160} \cdot 12^2 = \frac{43}{15} \simeq 2.87.$$

□

B.2 Estimating the L^1 distance to the Gaussian distribution

Let f_1 (resp. f_2) be the probability density of Y_1 (resp. Y_2), and let g be the Gaussian distribution with mean 0 and variance $1/12$. On Figure 4, we noticed that the probability densities f_1 and f_2 seem, on the average, to be very “close” to the Gaussian distribution g . The aim of the following proposition is to give upper-bounds for the L^1 distances $\|f_1 - g\|_{L^1} = \int_{\mathbb{R}} |f_1(x) - g(x)| dx$ and $\|f_2 - g\|_{L^1} = \int_{\mathbb{R}} |f_2(x) - g(x)| dx$.

Proposition 6: We have the following estimates :

$$\|f_1 - g\|_{L^1} \leq 0.07 \quad \text{and} \quad \|f_2 - g\|_{L^1} \leq 0.02.$$

The proof of Proposition 6 is given in Appendix. It combines exact inequalities and numerical estimates.

C. Posterior independence

In our study of the dequantized noise, we made the assumption that the X_k (resp. the $X_{k,l}$) are independent and uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$. Here we address the problem of the posterior independence of the dequantized noise, i.e. for example if we consider the dequantized noise Y_1 at two different points n and m ,

$$Y_1(n) = \sum_{k \in \mathbb{Z}} c_k X_{n+k} \quad \text{and} \quad Y_1(m) = \sum_{k \in \mathbb{Z}} c_k X_{m+k},$$

we are then interested in the correlation of $Y_1(n)$ and $Y_1(m)$. The result, which shows that we do not increase the correlation, is given by Proposition 7. We recall that the correlation coefficient of two random variables X and Y is defined by

$$\rho(X, Y) = \frac{E(XY) - E(X)E(Y)}{\sqrt{\text{var}(X)\text{var}(Y)}}.$$

For simplicity, the following proposition is stated and proved for the dequantized noise Y_1 . The result for Y_2 is the exact analogue.

Proposition 7: Let the X_k , for $k \in \mathbb{Z}$, be random variables uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$. Assume that for all $\delta \in \mathbb{Z}$, the correlation coefficient between X_k and $X_{k+\delta}$ is the same for all $k \in \mathbb{Z}$, and let us denote it by C_δ . Then, the correlation coefficient between $Y_1(n)$ and $Y_1(m)$ is $\rho(Y_1(n), Y_1(m)) = C_{m-n}$. In particular, this shows that if the X_k are independent, then the correlation coefficient of $Y_1(n)$ and $Y_1(m)$ is 0 when $n \neq m$.

Proof : Since, for all k , we have $E(X_k) = 0$, this implies that $E(Y_1(n)) = E(Y_1(m)) = 0$. We also have, for all $k \in \mathbb{Z}$, $\text{var}(X_k) = 1/12 = \text{var}(Y_1(n)) = \text{var}(Y_1(m))$. Thus, if we compute the ‘‘posterior’’ correlation coefficient, we get

$$\begin{aligned} \rho(Y_1(n), Y_1(m)) &= 12 \sum_{k \in \mathbb{Z}} \sum_{l \in \mathbb{Z}} c_k c_l E(X_{k+n} X_{l+m}) \\ &= \sum_{\delta \in \mathbb{Z}} \sum_{k \in \mathbb{Z}} c_k c_{k+\delta} C_{\delta+m-n} \end{aligned}$$

For $\delta \in \mathbb{Z}^*$, we have $\sum_{k \in \mathbb{Z}} c_k c_{k+\delta} = 0$. On the other hand, for $\delta = 0$, we already saw that $\sum_{k \in \mathbb{Z}} c_k^2 = 1$. Notice that these properties of the c_k are explained by the fact that the 1/2 Fourier translation is an isometry of L^2 . Finally, we obtain the announced result, which is $\rho(Y_1(n), Y_1(m)) = \sum_{k \in \mathbb{Z}} c_k^2 C_{m-n} = C_{m-n}$. \square

D. The flat regions model : final explanation of the dequantization effect

Let us summarize. We had written the “dequantized” signal as

$$S(n + \frac{1}{2}) = s(n + \frac{1}{2}) - \sum_{k \in \mathbb{Z}} X_{n+k} \frac{(-1)^k}{\pi(1/2 - k)}, \quad (8)$$

where $s(n + 1/2)$ is the original signal computed by Shannon interpolate at point $n + 1/2$, $S(n + 1/2)$ is the dequantized signal at the same point and $Y(n) = \sum_{k \in \mathbb{Z}} c_k X_{n+k}$ is the so-called “dequantized noise”. We have proven that $Y(n)$ is nearby Gaussian. Since in Proposition 1 we also prove that the addition of a Gaussian noise to the signal does not create any bias in the orientation map, we might be satisfied with this result. Now, we claim that the above explanation does not give an account of the change in the orientation histogram obtained by dequantization. Indeed, we prove in Proposition 4 that the addition to the signal of a uniformly distributed noise on an interval does not create a bias on the orientation larger than 4.7%. Thus something must be inaccurate in our assumptions. Actually, we notice that when the gradient of s is small at a given point, then the quantized values of $s(n)$ around this point are a very discrete signal. In other terms, assuming for simplicity that this point is 0 and that $s(0) = 0$, we have that $S(k) = s(k) - X_k \in \{0, 1, -1, 2, -2, ..\}$ where the first integer values are very majoritary. This means that $s(k)$ and X_k are highly correlated when the gradient is small. Thus, our model (8) explaining the good behaviour of $S(n + 1/2) = \text{original signal} + \sim \text{Gaussian noise} = s(n + \frac{1}{2}) + Y(n)$ will make sense only if we can point out that the dequantization process implies : $Y(n)$ and $s(n + 1/2)$ decorrelated. Now, using the same proof as in Proposition 7, we can show that this is not true. In fact, more precisely, by this result we have $\rho(Y(n), s(n + \frac{1}{2})) = \rho(X_n, s(n))$ under the sound assumption of stationarity. Thus, we gain or lose no independence of the signal and the noise obtained by dequantization. The final explanation will however arise from the technique developed above. We first point out (see Figures 5.c and 5.d) that all the bias in orientation histogram is due to low values of the quantized gradient, namely $|\nabla u| \leq 4$. The reason for this is the following : at a point (x, y) of an image where the gradient is large, the orientation is not much affected by the quantization. In fact, the angle error between the “true” orientation at the point and the orientation computed after the quantization of the image is proportional to $1/|\nabla u|$. Let us show this. We denote by u the original image, then $\nabla u = u_x + iu_y = |\nabla u| \exp(i\theta)$, where θ is the direction of the gradient. Let \tilde{u} denote the quantized image, and let q denote the quantization step. If $\tilde{\theta}$ denotes the direction of the gradient of the quantized image, we obtain

$\nabla \tilde{u} = |\nabla \tilde{u}| \exp(i\tilde{\theta}) = |\nabla u| \exp(i\theta) + z$, where z is a complex number (it represents the gradient of the difference between the true image and the quantized image) with modulus smaller than q . Thus, we get $|\sin(\theta - \tilde{\theta})| \leq q/|\nabla u|$. This shows that the points where the gradient is large are not much affected by the quantization effect.

The points with small gradient are majoritary (about 60%, see Figure 5.b) in the gradient norm histogram. Thus, we must focus on the points where $\frac{1}{2} \leq |\nabla u| \leq 4$. We notice that in a neighbourhood of such a point n , the histogram of values of $S(k)$ are a discrete uniform process centered at $S(n)$. Taking, without loss of generality, $S(n) = 0$, we can model the values around these points as discrete independent random values. See Figures 5.e and 5.f for the histogram and correlations. In flat regions, the gradient is quantized on a small number of values and we will see that the proposed Fourier translation has a strong dequantization effect. Let S denote the quantized signal. At a point n , we replace the quantized value $S(n)$, by the Shannon interpolate $S(n + 1/2)$, and then, we compute the gradient by

$$S(n + 1/2) - S(n - 1/2) = \sum_{k \in \mathbb{Z}} [S(n + k) - S(n + k - 1)] c_k.$$

In flat regions, we can assume that the difference $S(n) - S(n - 1)$ takes a small number of discrete values. For example, if we assume that it only takes the values 0, 1 or -1 , then the following proposition shows in particular that $S(n + 1/2) - S(n - 1/2)$ is no longer quantized.

Proposition 8: Let Z be the random variable defined by

$$Z = \sum_{k \in \mathbb{Z}} c_k Q_k,$$

where the Q_k are independent discrete random variables, taking the values 0, 1 or -1 , each one with probability $1/3$. Then Z follows the same probability distribution as

$$T_3 = \sum_{k \neq 2[3]} 3c_k X_k,$$

where the X_k are independent, uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$. Thus, for all $a < b$, $\mathbb{P}[a \leq Z \leq b] = \mathbb{P}[a \leq T_3 \leq b]$.

In fact, T_3 is nearby Gaussian. In particular, we have a nearby perfect dequantization since the previous proposition implies that for all $a < b$, then $\mathbb{P}[a \leq Z \leq b] > 0$.

The previous proposition can be extended to the case of a 2-D image in the following way. Let U denote the quantized image. Using the same notations as in Section 2, the gradient of U , before

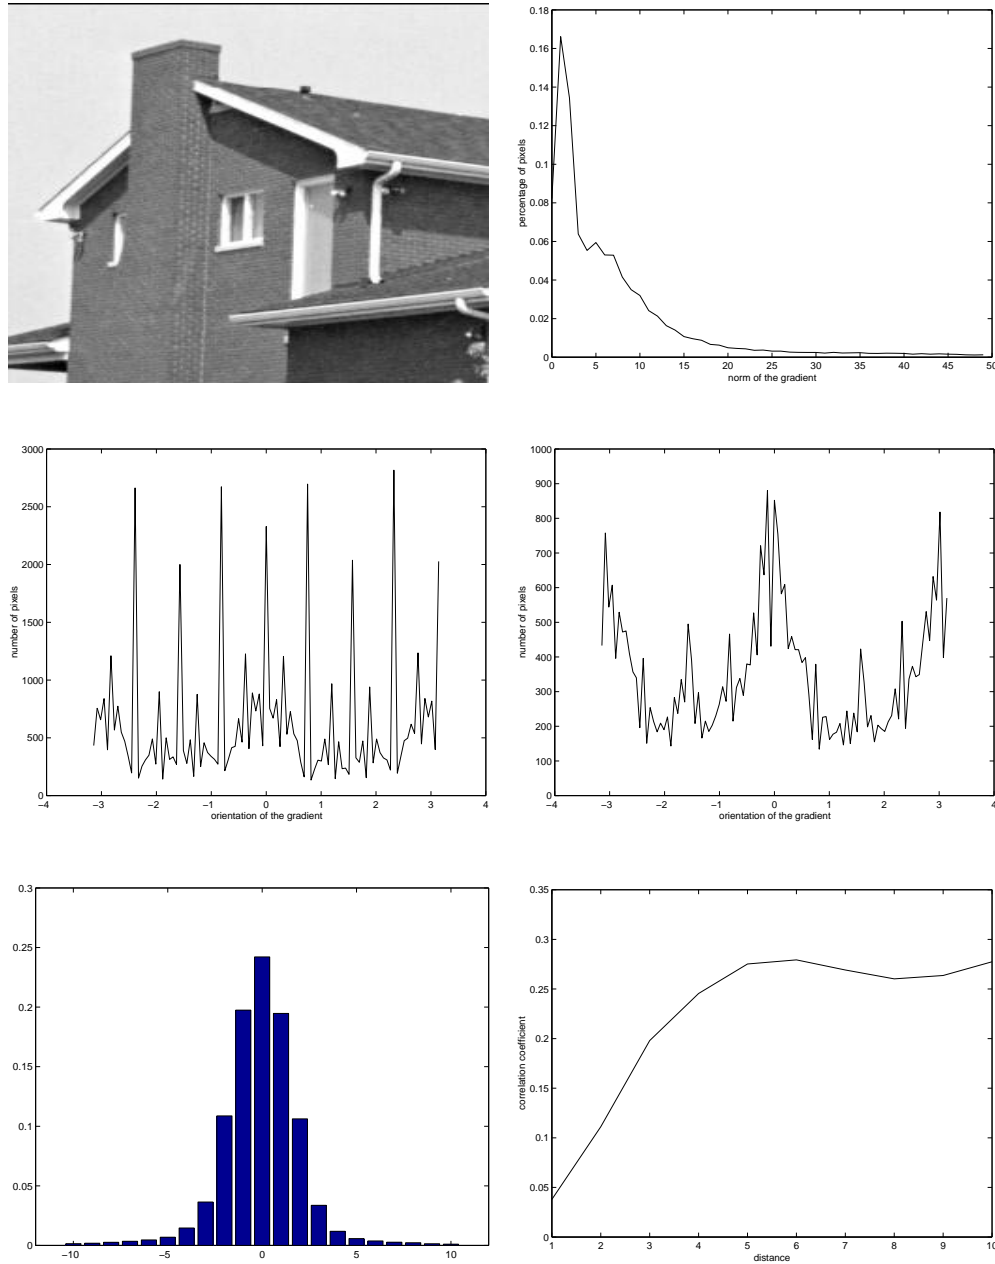


Fig. 5. Empirical observations on histograms and correlation. First row : left, the original image; right, histogram of the norm of the gradient. Second row : left, histogram of the gradient orientation for all points of the image; right, histogram of the gradient orientation for points with gradient norm larger than 5. Third row : left, local histogram (window of size 3×3) of $u(x) - u(x_0)$ for points x_0 such that $|\nabla u(x_0)| < 3$; right, correlation coefficient of $u(x_0 + d) - u(x_0)$ and $u(x_0 - d) - u(x_0)$ as a function of the distance d , for points x_0 such that $|\nabla u(x_0)| < 3$.

the 1/2-translation, has components A and B (in the referential defined by the (x, y) -axes rotated by $\pi/4$), where $A(n, m) = U(n + 1, m + 1) - U(n, m)$ and $B(n, m) = U(n, m + 1) - U(n + 1, m)$. After 1/2-Shannon translation, we obtain

$$A\left(n + \frac{1}{2}, m + \frac{1}{2}\right) = \sum_{k, l \in \mathbb{Z}} c_k c_l [U(n + k + 1, m + l + 1) - U(n + k, m + l)],$$

$$B\left(n + \frac{1}{2}, m + \frac{1}{2}\right) = \sum_{k, l \in \mathbb{Z}} c_k c_l [U(n + k, m + l + 1) - U(n + k + 1, m + l)].$$

Now, if we define for all $k, l \in \mathbb{Z}$, $Q_{k,l}^A = U(n + k + 1, m + l + 1) - U(n + k, m + l)$ and $Q_{k,l}^B = U(n + k, m + l + 1) - U(n + k + 1, m + l)$, we may assume that in flat regions all these variables are independent. Then, we can prove an analogous result as the one given by Proposition 8. More precisely, we can prove that A and B after translation have both the same nearby Gaussian distribution, which is the one of the random variable \tilde{T}_3 defined by

$$\tilde{T}_3 = \sum_{k, l \neq 2[3]} 3c_k c_l X_{k,l},$$

where the $X_{k,l}$ are independent, uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$.

IV. EXPERIMENTS AND APPLICATION TO THE DETECTION OF ALIGNMENTS

In this section, we present some applications of the proposed solution for dequantization. The first application is the detection of alignments in an image. In [3], we proposed a statistical criterion for the detection of meaningful alignments in an image. At each point of the image (with size $N \times N$), we compute an orientation $\Theta(x)$ which is orthogonal to the gradient at the considered point. Then, we consider a segment S in the image made of l points at distance 2 (thus, the gradients are computed on neighbourhoods that do not intersect, and we therefore make the assumption that they are independent). Let k be the number of points (among the l) which have their orientation aligned with the direction of the considered segment, at a given precision p (i.e. such that $|\Theta(x) - \Theta_0| \leq p\pi$, where Θ_0 denotes the orientation of the segment S). The probability of observing at least k such points on a length l segment is $P(k, l) = \sum_{j=k}^l \binom{l}{j} p^j (1-p)^{l-j}$ (because of the assumption that the orientations are independent and uniformly distributed on $[0, 2\pi]$). When this probability is very small, the event is highly non-casual and therefore meaningful. Generally, we compute meaningful alignments with the precision $p = 1/16$. But, sometimes, we are interested in alignments at a better precision, say for example $p = 1/64$. In Figure 6,

we first present the original image (upleft) : this is a result of the scan of Uccello's painting : "Presentazione della Vergine al tempio" (from the book *L'opera completa di Paolo Uccello*, Classici dell'arte, Rizzoli). This image is quantized on 32 grey levels. We first compute (upright) the meaningful alignments at precision $p = 1/16$. Then, we compute (downleft) the meaningful alignments at precision $p = 1/64$: it shows many diagonal alignments. These alignments are artefacts, their explanation is the quantization effect on the computation of orientations : directions multiple of $\pi/4$ are highly favoured. Last, we show the detection of meaningful alignments at precision $p = 1/64$ (downright), after the proposed solution for dequantization : $(1/2, 1/2)$ Fourier translation. The result shows that artefactual diagonal alignments are no longer detected.

We also noticed (but we have no theoretical argument to justify it) that the same method yields a significant improvement in orientation map of aliased images : see Figure 7. This is particularly true for aliasing due to direct undersampling, a barbaric but usual zoom-out method in many image processing softwares.

APPENDIX

Proof of Proposition 3 : Every linear combination of u_x and u_y is Gaussian because it is also a linear combination of the X_i , which are independent and Gaussian distributed. Thus (u_x, u_y) is a Gaussian vector. Since $\sum \lambda_i \mu_i = 0$, this implies that the correlation between u_x and u_y is 0. Since (u_x, u_y) is a Gaussian vector, this shows (see [6] for example) that u_x and u_y are independent. Moreover, the property $\sum \lambda_i^2 = \sum \mu_i^2$ shows that u_x and u_y are Gaussian with same mean and same variance. Finally, as in the proof of Proposition 1, the law of the couple (u_x, u_y) is given by a density function $f(x, y)$ which depends only on the radius $x^2 + y^2$ and not on the angle θ . Thus, θ is almost surely defined and uniformly distributed on $[0, 2\pi]$.

□

Proof of Proposition 4 : We use the same notations as in the proof of Proposition 1. The random variables $A = X_2 - X_3$ and $B = X_1 - X_4$ are independent and have the same density h , given by the convolution of the characteristic function of the interval $[-\frac{1}{2}, \frac{1}{2}]$ with itself, that is $h(x) = 1 - |x|$ for $|x| \leq 1$, and $h(x) = 0$ otherwise. Now we compute the law of $\alpha = \theta - \pi/4$, knowing from (3) that $B = A \tan(\alpha)$. Thanks to symmetries, we first consider the case $0 \leq \alpha \leq \pi/4$. The distribution function of α is $F(\alpha) = P[0 \leq B \leq A \tan \alpha]$, that is

$$F(\alpha) = \int_{x=0}^1 \left(\int_{y=0}^{x \tan \alpha} (1-y) dy \right) (1-x) dx.$$

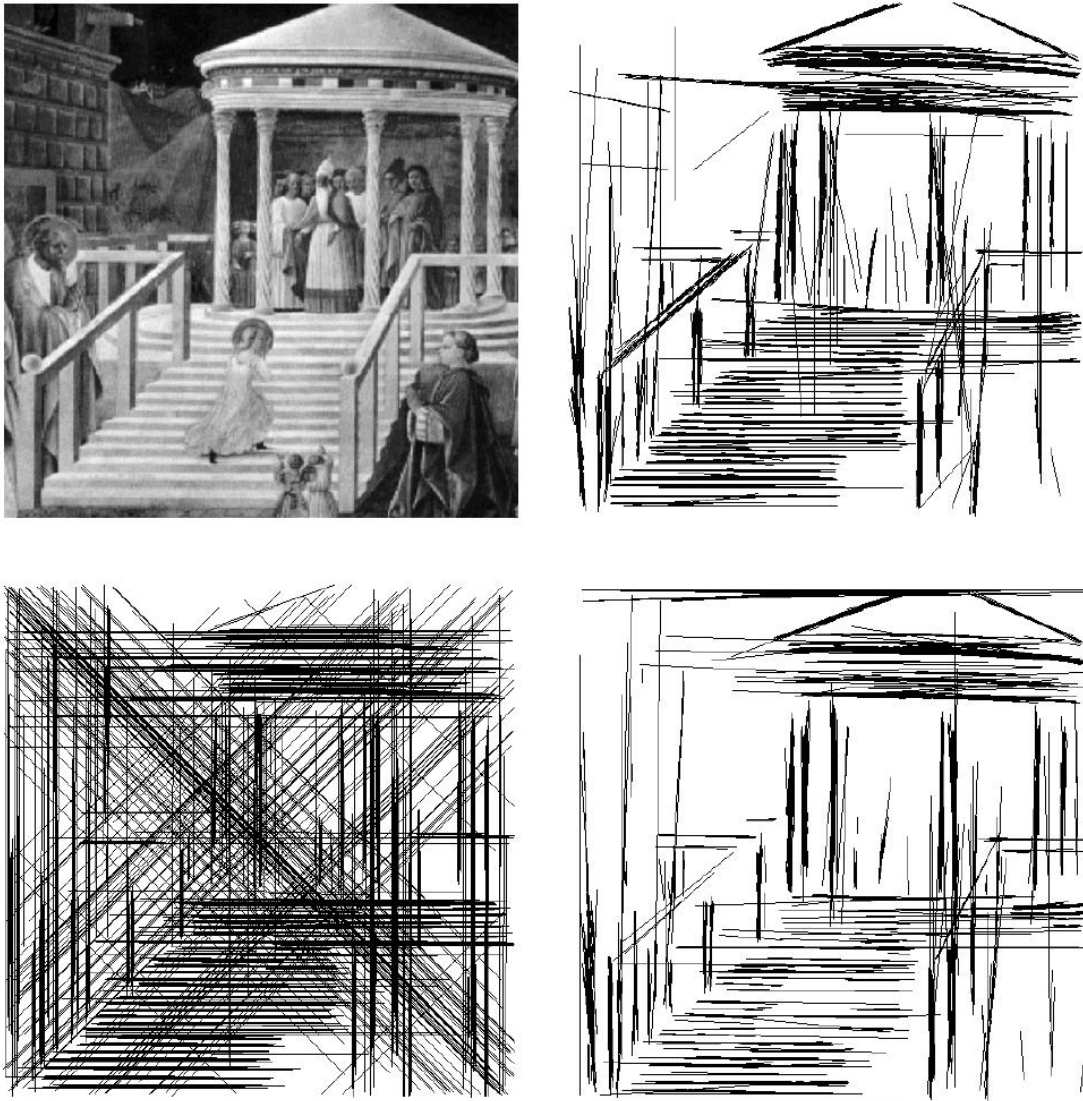


Fig. 6. The effect of quantization on the detection of alignments. First row : left, the original painting image, quantized on 32 grey levels, right, the meaningful alignments for precision $p = 1/16$. Second row : left, the meaningful alignments for precision $p = 1/64$; right, the meaningful alignments for precision $p = 1/64$, after $(1/2, 1/2)$ -translation.




Fig. 7. The effect of the Fourier translation on an aliased image. First row : left, the aliased image; right, the image after $(1/2, 1/2)$ Fourier translation. Second row : left, the histogram of the gradient orientation in the aliased image; right, the histogram of the gradient orientation after translation.

Hence, the law of $\alpha \in [0, \pi/4]$ is given by the density function

$$f(\alpha) = F'(\alpha) = \int_0^1 x(1 + \tan^2 \alpha)(1 - x \tan \alpha)(1 - x) dx = \frac{1}{12}(1 + \tan^2 \alpha)(2 - \tan \alpha).$$

Finally, since $\alpha = \theta - \pi/4$ and by symmetries, we obtain the announced law for θ .

□

Proof of Proposition 6 : The c_k are symmetric around $1/2$, i.e. $c_{1-k} = c_k$. Thus, we can write $Y = \sum_{k \geq 1} c_k(X_k + X_{1-k})$. Let us denote

$$Z_n = \sum_{k=1}^n c_k(X_k + X_{1-k}),$$

and let f_n be the probability density of Z_n . We then have $f_n = f_{n-1} * h_n$, where h_n is the probability density of $c_n(X_n + X_{1-n})$. We first prove the following proposition

Proposition 9: For all $n \geq 1$, we have

$$\|f - g\|_{L^1} \leq \|f_n - g\|_{L^1} + \frac{1}{12} \|g''\|_{L^1} \sum_{k \geq n+1} c_k^2.$$

Proof: For $n \geq 2$, we have $f_n = f_{n-1} * h_n$, where h_n is the probability density of $c_n(X_n + X_{1-n})$. We notice that h_n is a positive even function, with compact support $[-|c_n|, |c_n|]$ and satisfying $\int h_n = 1$. Thus,

$$\begin{aligned} \|f_n - g\|_{L^1} &\leq \|f_{n-1} * h_n - g * h_n\|_{L^1} + \|g * h_n - g\|_{L^1} \\ &\leq \|f_{n-1} - g\|_{L^1} + \|g * h_n - g\|_{L^1}. \end{aligned}$$

We now compute $\|g * h_n - g\|_{L^1}$. For $x \in \mathbb{R}$, using the definition of $g * h_n$ and the integral Taylor formula, we get, $yh_n(y)$ being odd,

$$g * h_n(x) - g(x) = \int_{-|c_n|}^{|c_n|} y^2 \left(\int_0^1 (1-t) g''(x+ty) dt \right) h_n(y) dy.$$

Then, we can estimate the L^1 distance, and obtain

$$\|g * h_n - g\|_{L^1} \leq \frac{1}{12} \|g''\|_{L^1} c_n^2.$$

We add these inequalities and thus obtain the announced result. \square

Using the previous Proposition, in order to have a numerical estimate of $\|f - g\|$, we can use a computational software to compute numerically the first terms $f_1, f_2, \dots, f_{10}, \dots$ and on the other hand compute an upper-bound for the tail $\|g''\| \sum_{k \geq N+1} c_k^2$. We first compute $\|g''\|_{L^1}$. For $x \in \mathbb{R}$, we have $g''(x) = (x^2/\sigma^4 - 1/\sigma^2)g(x)$, where $\sigma^2 = 1/12$ is the variance of g . Thus, using an integration by parts and the properties $\int_{\mathbb{R}} g = 1$ and $\int_{\mathbb{R}} x^2 g(x) = \sigma^2$, we get $\|g''\| = 4g(\sigma)/\sigma \leq 12$. Then, we compute the rest of the sum $\sum c_k^2$, using a comparison with an integral, and get $\sum_{k \geq N+1} c_k^2 \leq 1/\pi^2 N$. Thus, we obtain

$$\frac{1}{12} \|g''\|_{L^1} \sum_{k \geq N+1} c_k^2 \leq \frac{1}{\pi^2 N}.$$

We can now compute an upper-bound for the L^1 distance between the density function of Y and the Gaussian distribution. We can also do this in a similar way for the random variable Y_2 (dequantized noise in dimension 2). The numerical estimates given in Proposition 6 are obtained using Proposition 9 for $n = 25$, and the numerical computation of $\|f_{25} - g\|_{L^1}$. \square

Proof of Proposition 8 : Let $\delta(x)$ denote the dirac function. Then, the probability distribution of the Q_k is $d(x) = 1/3 \cdot (\delta(x - 1) + \delta(x) + \delta(x + 1))$. The main point is to notice that the result of the convolution of $d(x)$ with the uniform distribution on $[-\frac{1}{2}, \frac{1}{2}]$, is the uniform distribution on $[-\frac{3}{2}, \frac{3}{2}]$. This means that $Q_k + X_k$ has the same probability distribution as $3X_k$. And consequently, $\sum_k c_k(Q_k + X_k)$ has the same probability distribution as $\sum_k 3X_k$. We now consider the Fourier transform of the previous distributions (it will convert the convolution into a product). Let $F_1(t)$ denote the Fourier transform of $\sum_k c_k X_k$. From (7), we know that

$$F_1(t) = \prod_{k \in \mathbb{Z}} \operatorname{sinc} \left(\frac{c_k t}{2} \right) = \prod_{k \in \mathbb{Z}} \operatorname{sinc} \left(\frac{t}{\pi(2k - 1)} \right).$$

We denote by $F_3(t)$ the Fourier transform of $\sum_k 3c_k X_k$. We then have $F_3(t) = F_1(3t)$. On the other hand, the Fourier transform of $d(x)$ is $(1 + 2 \cos t)/3$. Thus, if we denote by $G(t)$ the Fourier transform of the law of $\sum_k c_k Q_k$, we have

$$G(t) = \prod_{k \in \mathbb{Z}} \left(\frac{1 + 2 \cos(c_k t)}{3} \right),$$

where in both cases the convergence of products is uniform on every compact subset of \mathbb{R} . Since $\sum_k c_k(Q_k + X_k)$ has the same probability distribution as $\sum_k 3X_k$, this shows that $G(t)F_1(t) = F_3(t)$ for any real t . We now show that there exists a continuous function H_1 such that for all t , $F_3(t) = F_1(t)H_1(t)$. In fact, we have $F_3(t) = \prod_{k \in \mathbb{Z}} \operatorname{sinc} \left(\frac{3t}{\pi(2k-1)} \right)$, and we decompose this product into two products : the first one is the product over all k such that $k \neq 2[3]$, and the second is for all k such that $k = 2[3]$. Now, for $k = 2[3]$, we can write $k = 3k' - 1$, and thus $3/(2k - 1) = 1/(2k' - 1)$. This shows that

$$\prod_{k=2[3]} \operatorname{sinc} \left(\frac{3t}{\pi(2k - 1)} \right) = F_1(t),$$

and consequently, we have $F_3(t) = F_1(t)H_1(t)$, where

$$H_1(t) = \prod_{k \neq 2[3]} \operatorname{sinc} \left(\frac{3t}{\pi(2k - 1)} \right),$$

and it may be shown that H_1 is continuous on \mathbb{R} . Thus, we have for all t , $G(t)F_1(t) = H_1(t)F_1(t)$. Since the zeros of F_1 are discrete, and G and H_1 are continuous, this shows that

$$\forall t \in \mathbb{R}, \quad G(t) = H_1(t) = \prod_{k \neq 2[3]} \operatorname{sinc} \left(\frac{3t}{\pi(2k - 1)} \right).$$

And thus, $Z = \sum_k c_k Q_k$ has the same probability distribution as $\sum_{k \neq 2[3]} 3c_k X_k$. Moreover, thanks to Levy Theorem, we have a convergence in law of the partial sums $\sum_{|k| \leq n} c_k Q_k$ to T_3 .

□

Acknowledgements :

This work was supported by Office of Naval Research under grant N00014-97-1-0839. We thank Richard Lau for his personal interest and encouragement. We thank Bernard Rougé for his valuable suggestions and comments.

REFERENCES

- [1] L.Alvarez, F.Guichard, P.L.Lions and J.M.Morel, "Axioms and fundamental equations of image processing : multiscale analysis and PDE", *Arch. for Rat. Mech. (16)*, No. 9, pp. 200-257, 1993.
- [2] A.Canny, "A computational approach to edge detection", *IEEE Trans. on PAMI (8)*, No. 6, pp. 679-698, 1986.
- [3] A.Desolneux, L.Moisan and J.-M.Morel, "Meaningful Alignments", *IJCV(40)*, No. 1, pp. 7-23, 2000.
- [4] J. Dieudonné, *Calcul infinitésimal*, 1980, Ed. Hermann.
- [5] P. Pirsch and A.N. Netravali, "Transmission of Gray Level Images by Multilevel Dither Techniques", *Computers and Graphics (7)*, pp. 31-44, 1983.
- [6] W. Feller, *An introduction to probability theory and its applications*, vol. II, 1971, Wiley Series in Probability and Mathematical Statistics.
- [7] J. Lopez-Krahe and P. Pousset, "Transformée de Hough discrète et bornée, application à la détection de droites parallèles et du réseau routier", *Traitement du Signal (5)*, No. 4, pp. 281-290, 1988.
- [8] E. Lutton, H. Maître and J. Lopez-Krahe, "Contribution to the Determination of Vanishing Points Using Hough Transform. ", *IEEE Trans. on PAMI (16)*, No. 4, pp. 430-438, 1994.
- [9] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, 1998.
- [10] D.Marr and E.Hildreth, "Theory of Edge Detection", *Proc. Roy. Soc. Lond.*, B207, pp. 187-217, 1980.
- [11] D. Mumford, *Pattern Theory and Vision*, Notes de cours du centre Emile Borel, Questions mathématiques en traitement du signal et de l'image (14 sept.98 - 18 dec.98), chapt. 5 (No 16).
- [12] P. Perona, "Orientation Diffusions", *IEEE Trans. on Image Processing (7)*, No. 3, pp. 457-467, 1998.
- [13] G.Sapiro and A.Tannenbaum, "Affine Invariant Scale-Space", *IJCV (11)*, No. 1, pp. 25-44, 1993.
- [14] B.Tang, G.Sapiro and V. Caselles, "Direction Diffusion", *ICCV*, pp. 1245-1252, 1999.
- [15] A.P.Witkin, "Scale-space Filtering", *IJCAI*, pp. 1019-1022, 1983.