Int J Comput Vis
DOI 10.1007/s11263-009-0310-5

Study of a Robust Feature: The Pointwise Lipschitz Regularity

C. Damerval · S. Meignen

Received: 13 August 2008 / Accepted: 5 November 2009

Abstract The aim of this paper is to highlight the relevance in computer vision of the pointwise Lipschitz regularity $\alpha \in \mathbb{R}$. The regularity α gives a measure of the local regularity of the intensity function associated to an image. Known wavelet methods provide an efficient computation of α at contour points of the image. From a theoretical point of view, we study the effect of geometric deformations and other specific transformations applied to the image, showing invariance properties. From a practical point of view, we assess the robustness of the regularity α when the image undergoes various transformations. The results we obtain show the Lipschitz regularity α is a suitable feature for applications in computer vision.

Keywords Lipschitz regularity · Invariance properties · Multiscale contour point detector · Extraction of characteristic values · Robustness to geometric deformations and transformations applied to the image

1 Introduction

The extraction of invariant or robust features from an image appears as a central issue in computer vision. The difficulty of this problem lies in the fact that natural scenes are often viewed under various situations, corresponding to a wide class of transformations (geometric deformations or illumination change for instance). So as to reach certain invariance properties regarding a transformation (scale change for instance), one should either extract features that are invariant (direct approach), or take into account the transformation (indirect approach). As examples of indirect approaches, let us cite methods based on *Scale-Space* theory (Iijima 1962; Lindeberg 1994; Witkin 1983). These evidence regions of interest, which are stable through local geometric deformations (Lindeberg and Garding 1997; Lindeberg 1998). These regions are identified by their location and characteristic scale, and their content can be quantified by a suitable descriptor which makes up a robust feature (Lowe 2004). Recent works compared state-of-the-art interest regions detectors (Mikolajczyk et al. 2005) and region descriptors (Mikolajczyk and Schmid 2005). Existing methods proved to be efficient for one type of scene or transformation (Baumberg 2000; Matas et al. 2002; Kadir et al. 2004; Mikolajczyk and Schmid 2004; Tuytelaars and Gool 2004). However, no method outperforms the others in all cases, so combining different kinds of features seems relevant. Let us also mention the large work done on the extraction of local features based on differential invariants (Koenderink and Doorn 1987; Florack et al. 1993). These methods are efficient in practice, but their theoretical invariance properties are not always known, and they often entail complex computations (Ko and Lee 2008; Brown and Lowe 2007).

In this paper, we study a feature related to the local regularity of the intensity function: the pointwise Lipschitz regularity $\alpha \in \mathbb{R}$ (denoted regularity α). In a 1D context, the regularity α was mostly studied in the case of multifractal signals (Arneodo et al. 1993, 1997, 1998, 1999; Benassi et al. 1998). Successful applications to the characterization of

C. Damerval (✉)
Dept. of Computer Science, K.U. Leuven, Celestijnenlaan 200A, Postbus 2402, 3001 Heverlee, Belgium
e-mail: christophe.damerval@cs.kuleuven.be

S. Meignen
Laboratory Jean Kuntzmann, University of Grenoble, LJK Site Campus–Tour IRMA, 51 rue des Mathmatiques, 38400 Saint Martin d'Hères, France
e-mail: sylvain.meignen@imag.fr

Published online: 24 November 2009

singularities (Jaffard and Meyer 1996; Mallat and Hwang 1992) and landmark registration (Bigot 2003) were also carried out. In a 2D context, methods based on regularity measures were also put forward for the analysis, the synthesis and the classification of textured images (Deguy et al. 2000; Kaplan and Kuo 1995a, 1995b). Note that in these applications, the regularity is used from a global point of view, often assuming a stationary hypothesis. Here we focus on the pointwise regularity α. As we will see, the regularity α appears as a relevant feature: on the one hand there exist some theoretical invariance properties (especially regarding geometric transformations), and on the other hand values of regularity α can be efficiently computed. This framework leads to a direct approach for extracting a pointwise feature: the regularity α associated to contour points.

This paper is organized as follows. In a first part we investigate invariance properties of the regularity α. We show some theoretical results in 1D on α-Lipschitz functions, laying the emphasis on geometric deformations. We also present invariance properties of the regularity α in 2D. In a second part, we put forward a direct approach of feature extraction, based on a robust computation of the regularity α. So as to numerically estimate α in 2D, we recall an algorithm using a multiscale contour detector (Mallat and Zhong 1992). This gives pointwise estimations of α at contour points of the image. We show the robustness of this estimation firstly in controlled situations (noise addition, deformations applied to geometric images), and secondly in the case of a known affine deformation. In a third part, we study the robustness of our approach for a wide class of image transformations (geometric and photometric). In particular we design an evaluation procedure and apply it to natural scenes viewed under various imaging conditions. Results show that the regularity α makes up a robust feature.

2 Study in the Monodimensional Case

Here we recall the definition of the Lipschitz regularity of a monodimensional real function, and we investigate the impact of certain transformations on an α-Lipschitz function. We quote the main results (whose proofs are detailed in Appendix A). These provide sufficient conditions for the preservation of the Lipschitz regularity.

Definition 1 (Pointwise Lipschitz regularity) A function $f : \mathbb{R} \to \mathbb{R}$ is α-Lipschitz at $x_0 \in \mathbb{R}$

- for $\alpha \in]0, 1[$, if there exists a neighborhood V of x_0 and $A > 0$ such that

$$\forall x \in V, \quad |f(x) - f(x_0)| \leq A|x - x_0|^\alpha \tag{1}$$

- for $\alpha > 1$ (α non integer), denoting $n = \lfloor \alpha \rfloor$, if there exists a neighborhood V of x_0, $A > 0$ and a polynomial $P_n(x)$ of order n ($n \leq \alpha < n + 1$, P_n depending on x_0) such that

$$\forall x \in V, \quad |f(x) - P_n(x)| \leq A|x - x_0|^\alpha \tag{2}$$

- for $\alpha = n \in \mathbb{N}^*$, if f is C^n at x_0.

The point of studying the pointwise Lipschitz regularity lies in the fact many functions contain singularities. For instance, the real function $x \mapsto \sqrt{|x|}$ is C^∞ at any $x \neq 0$ ($x \in \mathbb{R}$), and α-Lipschitz at $x = 0$, for $\alpha \leq 1/2$.

Remark If $f : \mathbb{R} \to \mathbb{R}$ is α-Lipschitz at any $x_0 \in \mathbb{R}$, f is said uniformly α-Lipschitz.

2.1 Case $\alpha > 0$

Proposition 1 *Let $x_0 \in \mathbb{R}$ and $f, g : \mathbb{R} \to \mathbb{R}$ related by $g(x) = f(u^{-1}(x))$ where $u : \mathbb{R} \to \mathbb{R}$. We assume that u is C^∞, invertible and linear (so that u^{-1} is linear)*

$$\forall \alpha > 0, \quad f \text{ } \alpha\text{-Lipschitz at } x_0 \implies g \text{ } \alpha\text{-Lipschitz at } x_0 \tag{3}$$

Proposition 2 *Let $x_0 \in \mathbb{R}$ and $f, g : \mathbb{R} \to \mathbb{R}$ related by $g(x) = c(x)f(x) + d(x)$ where $c, d : \mathbb{R} \to \mathbb{R}$. In the case $\alpha \in]0, 1[$ and c, d are C^1, or in the case $\alpha > 1$ and c, d are C^∞:*

$$\forall \alpha > 0, \quad f \text{ } \alpha\text{-Lipschitz at } x_0 \implies g \text{ } \alpha\text{-Lipschitz at } x_0 \tag{4}$$

2.2 Case $\alpha \leq 0$

In practice there are often non-differentiable cases, thus leading to negative values of α. In this situation, the distribution theory is useful (Rudin 1991), in particular we recall that tempered distributions are defined on

$$\mathcal{S}(\mathbb{R}) = \left\{ \begin{array}{l} \varphi : \mathbb{R} \to \mathbb{R}, C^\infty, \text{fast decaying} : \forall n \in \mathbb{N}, \\ \exists C > 0, \exists N \in \mathbb{N}^*, \|\varphi^{(n)}\|_\infty \leq \frac{C}{(1+x^2)^N} \end{array} \right\}$$

Here we consider tempered distributions of finite order, i.e., for which there exists a function $f : \mathbb{R} \to \mathbb{R}$ so that, denoting $f^{(n)}$ the n-th derivative in the sense of distributions,

$$\forall \varphi \in \mathcal{S}(\mathbb{R}), \quad \langle T, \varphi \rangle = \langle f^{(n)}, \varphi \rangle \tag{5}$$

Now, let us investigate the influence of certain transformations on an α-Lipschitz distribution.

Definition 2 Let T a tempered distribution of finite order. T is said uniformly α-Lipschitz if its primitive is uniformly $\alpha+1$-Lipschitz.

Definition 3 Given a distribution T, the distribution T^u associated to the change of variables $y = u^{-1}(x)$ (assuming $u : \mathbb{R} \to \mathbb{R}$ is C^∞, invertible and such that $\varphi \in \mathcal{S}(\mathbb{R}) \Rightarrow (\varphi \circ u)u' \in \mathcal{S}(\mathbb{R})$ is given by

$$\langle T^u, \varphi \rangle = \langle T, (\varphi \circ u)u' \rangle \tag{6}$$

Remark This definition comes from the fact that if two functions f and g are related by $g(x) = f(u^{-1}(x))$, their associated distributions T_f and T_g are related by

$$\begin{aligned}\langle T_g, \varphi \rangle &= \int_{\mathbb{R}} f(u^{-1}(x))\varphi(x)dx \\ &= \int_{\mathbb{R}} f(y)\varphi(u(y))u'(y)dy \\ &= \langle T_f, (\varphi \circ u)u' \rangle\end{aligned}$$

Proposition 3 *Let T a distribution of finite order and T^u the associated distribution according to Definition 3. For all $\alpha \leq 0$,*

$$T\ \alpha\text{-Lipschitz} \implies T^u\ \alpha\text{-Lipschitz} \tag{7}$$

Proposition 4 *Let T a distribution of finite order, and T_1 the distribution given by $T_1 = cT + d$, assuming $c, d : \mathbb{R} \to \mathbb{R}$ are bounded and in $\mathcal{S}(\mathbb{R})$ (so that T_1 is well defined on $\mathcal{S}(\mathbb{R})$). For all $\alpha \leq 0$,*

$$T\ \alpha\text{-Lipschitz} \implies T_1\ \alpha\text{-Lipschitz} \tag{8}$$

Conclusion The presented results show that when a 1D function undergoes regular geometric transformations, or a change in contrast and/or illumination, its Lipschitz regularity is not altered (under certain sufficient conditions).

3 Study in the Bidimensional Case

From here, we consider an image given by its intensity function $f : \mathbb{R}^2 \to \mathbb{R}$. We focus on the pointwise Lipschitz regularity of f (denoted regularity α). First we present the notion of regularity α in 2D, as a generalization of the 1D case. Then we investigate its invariance properties, especially in the case of geometric transformations.

3.1 Notion of Regularity α in 2D

Definition 4 Let $f : \mathbb{R}^2 \to \mathbb{R}$ and $x_0 \in \mathbb{R}^2$. Given $\theta \in [0, \pi[$, we define $f_\theta : \mathbb{R} \to \mathbb{R}$ as $f_\theta(h) = f(x_0 + hu_\theta)$, where $u_\theta = (\cos\theta, \sin\theta)$. For $\alpha \in \mathbb{R}$, f is α-Lipschitz at $x_0 \in \mathbb{R}^2$ if

$$\forall \theta \in [0, \pi[, \quad f_\theta\ \alpha\text{-Lipschitz at } 0 \tag{9}$$

Remark When $\alpha \in]0, 1[$, provided f is α-Lipschitz at x_0, we can write (for x in a neighborhood of x_0):

$$\begin{aligned}|f(x) - f(x_0)| &= |f(x_0 + (h\cos\theta, h\sin\theta)) - f(x_0)| \\ &= |f_\theta(h) - f_\theta(0)| \\ &\leq Ah^\alpha \\ &\leq A\|x - x_0\|^\alpha\end{aligned}$$

So Definition 5 agrees with the usual definition of the Lipschitz regularity and also fits with the case $\alpha \leq 0$.

Definition 5 (Regularity α) Let $f : \mathbb{R}^2 \to \mathbb{R}$ and $x_0 \in \mathbb{R}^2$. The regularity α of f at x_0 is defined as

$$\alpha = \inf\{\alpha_0 \in \mathbb{R},\ f\ \alpha_0\text{-Lipschitz at } x_0\} \tag{10}$$

We represent on Fig. 1 common image features and the corresponding values of the regularity α. Note that in the case of an image representing a contour (a line for example) f is regular along the tangent of the contour and irregular along the normal direction (see Fig. 2). More generally, so as to estimate precisely the regularity α of f at a given point, it is important to determine the direction of maximum irregularity. We explain further how to compute this direction and estimate α (see Sect. 4).

3.2 Invariance Properties

Here we present results on the preservation of the regularity α in 2D when the image undergoes transformations. First, let

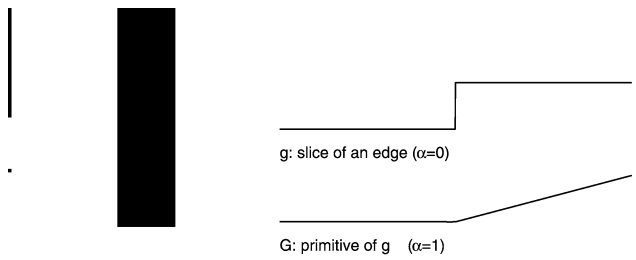

Fig. 1 (**a**) Regularity α associated to common image features

$$\begin{cases}\text{Edge}, & \alpha = 0 \\ \text{Line}, & \alpha = -1 \\ \text{Isolated point}, & \alpha = -2\end{cases}$$

(**b**) Illustration for an edge (see (**a**)). An horizontal slice (direction of maximum irregularity) leads to a 1D function g. The regularity of its primitive G is $\alpha_G = \alpha_g + 1$. Since G is piecewise linear, $\alpha_G = 1$, and so $\alpha_g = 0$. Thus the regularity α of an edge equals 0

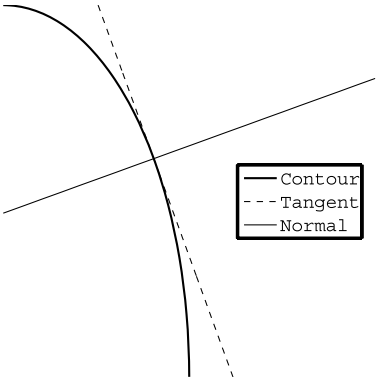

Fig. 2 At a point belonging to a contour line, the Lipschitz regularity is minimal along the normal direction; along this direction, the regularity α can be accurately computed

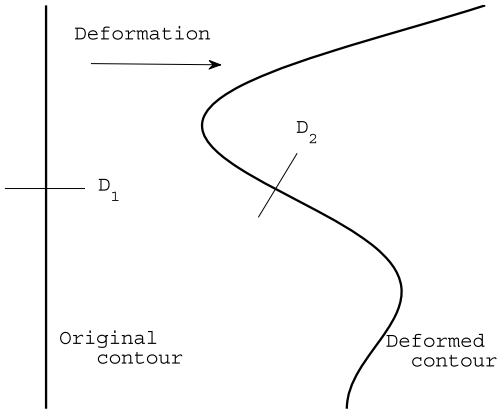

Fig. 3 If a contour undergoes a deformation which does not change its topology, the regularity α is preserved (D_1, D_2: directions of maximum irregularity)

us study the case of a contrast and/or illumination change, formulated as: $\forall x \in \mathbb{R}^2$,

$$g(x) = cf(x) + d, \quad \text{with } c \neq 0 \text{ and } d \in \mathbb{R} \tag{11}$$

We denote $\alpha \in \mathbb{R}$ (resp. α') the pointwise regularity of f (resp. of g) at x_0. We can write

$$\alpha' = \inf\{\alpha_0 \in \mathbb{R},\ g\ \alpha_0\text{-Lipschitz at } x_0\}$$

$$\alpha' = \inf\{\alpha_0 \in \mathbb{R}, \forall \theta \in [0, \pi[,\ g_\theta\ \alpha_0\text{-Lipschitz at } 0\}$$

$$\text{with } g_\theta = g(x_0 + hu_\theta) = cf(x_0 + hu_\theta) + d = cf_\theta + d$$

Since $g_\theta = cf_\theta + d$ and $f_\theta = \frac{1}{c}g_\theta - \frac{d}{c}$, the monodimensional functions g_θ and f_θ have the same Lipschitz regularity (according to Proposition 2). This is valid for all $\theta \in [0, \pi[$, and therefore f and g have the same regularity $\alpha = \alpha'$. So the regularity α is invariant in the case of a constant change of contrast or illumination.

Secondly, let us study the case of a constant affine deformation (so including rotation and scale change), widely studied in Scale-Space theory (Lindeberg 1994). This can be formulated as: $\forall x \in \mathbb{R}^2$,

$$g(x) = f(Bx), \quad \text{with } B \text{ a } 2 \times 2 \text{ invertible matrix} \tag{12}$$

We assume that the pointwise regularity of f at $x_0 \in \mathbb{R}^2$ is $\alpha \in \mathbb{R}$. According to Definition 5, there exists $A > 0$ so that

$$\forall \theta \in [0, \pi[, \quad |f(x_0) - f(x_0 + hu_\theta)| \leq Ah^\alpha \tag{13}$$

Let us study the Lipschitz regularity of g at $y_0 = B^{-1}x_0$. For $\theta \in [0, \pi[$ we have

$$\begin{aligned} |g(y_0) - g(y_0 + hu_\theta)| &= |g(B^{-1}x_0) - g(B^{-1}x_0 + hu_\theta)| \\ &= |f(x_0) - f(x_0 + hBu_\theta)| \end{aligned}$$

Since $Bu_\theta = \lambda u_{\theta'}$ with $\lambda \in \mathbb{R}$ ($\lambda \neq 0$) and $\theta' \in [0, \pi[$:

$$\begin{aligned} |g(y_0) - g(y_0 + hu_\theta)| &= |f(x_0) - f(x_0 + h\lambda u_{\theta'})| \\ &\leq (A|\lambda|^\alpha)h^\alpha \end{aligned}$$

So there exists $A' > 0$ so that

$$\forall \theta \in [0, \pi[, \quad |g(y_0) - g(y_0 + hu_\theta)| \leq A'h^\alpha \tag{14}$$

and g is α-Lipschitz at x_0.

Now, let us assume the regularity α of f corresponds to a minimum α_0 attained in a certain direction θ_0. Since we consider a constant affine deformation, there exists θ_1 so that Bu_{θ_1} and u_{θ_0} are collinear. Hence, the regularity α of g at $B^{-1}x_0$ corresponds to a minimum α_0 in the direction θ_1. So the regularity α is preserved when a constant affine deformation is applied to the image.

Finally, let us discuss more general transformations. Note that the 1D results shown in Sect. 2 regarding the preservation of the Lipschitz regularity only provide sufficient conditions (which may be not necessary). In this regard, it is important to note that when $g(x) = f(v(x))$, $x \in \mathbb{R}^2$ (with $v : \mathbb{R}^2 \to \mathbb{R}^2$), the regularity α is not necessarily preserved: provided v is regular enough, g may be more regular than f, resulting in a higher regularity α for g than for f. Nevertheless note that in most applications, the topology of the contours is not altered by transformations applied to the image (like local affine deformations). In these cases, the regularity α on these contours is preserved (see Fig. 3).

4 Robust Computation of the Regularity α

We explain here how to compute values of α ($\alpha \in \mathbb{R}$) at contour points (see Fig. 4). Such points appear as relevant, since they correspond to singularities in the image, where α can be

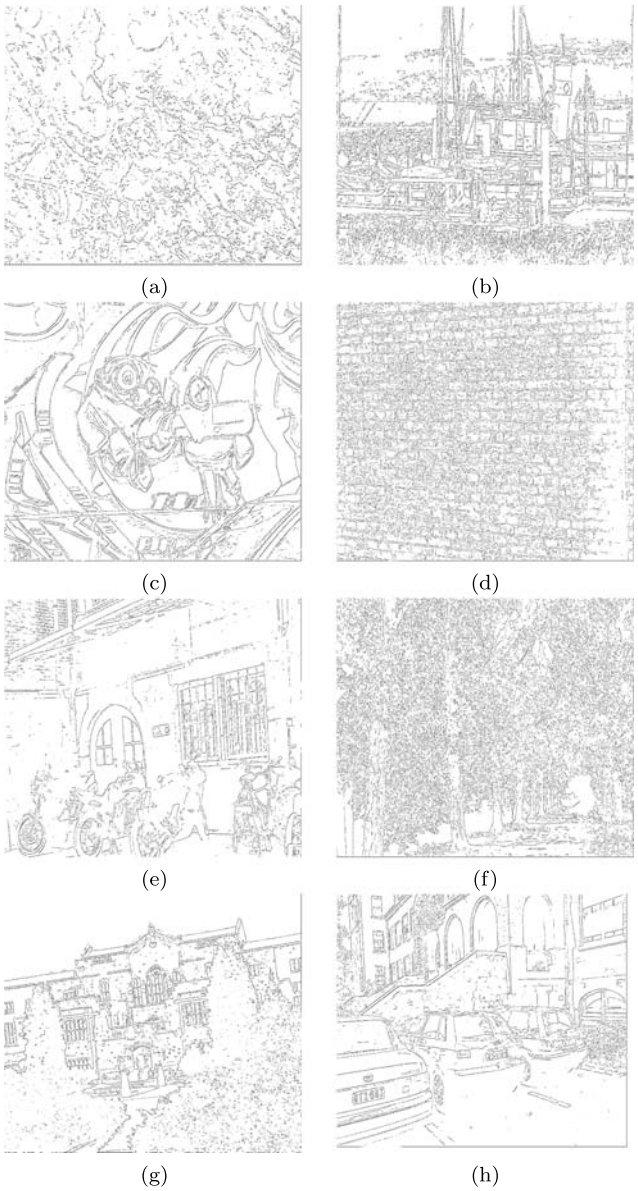

Fig. 4 Examples of detected contour points, associated to images of the data set used in Sect. 5 (more particularly, (**a**)–(**h**) correspond to the images on Figs. 12(a)–(h), left part)

accurately computed; besides, the kind of contour (an edge for example) can also be characterized by the regularity α.

Mainly the computation of α can be achieved thanks to two facts: on the one hand, there is a link between wavelets and Lipschitz regularity (see Appendix B); on the other hand, there exist efficient algorithms that implement wavelet transforms.

4.1 Wavelet Multiscale Detector of Contour Points

We recall here a known method which allows to detect contour points and compute values of regularity α at all these points. Let us recall that a maximum in the sense of Canny

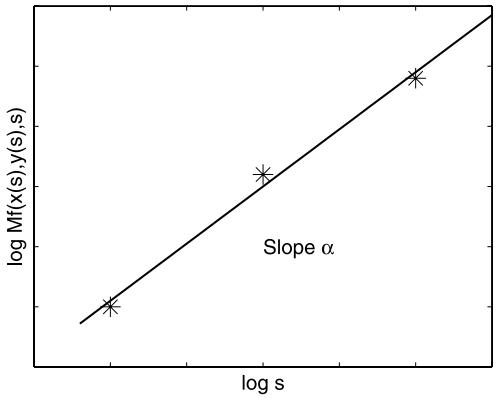

Fig. 5 Estimation of the regularity α: perform a linear regression of $\log Mf$ (modulus of the gradient) over $\log s$ (scale). The regularity α is given by the slope of the regression (3 fine scales are sufficient for an accurate estimation)

(MC) is a point $(x, y) \in \mathbb{R}^2$ where the magnitude of the gradient attains a local maximum in the direction of the gradient (Canny 1986). Besides, a multiscale formulation of Canny's detector was introduced in Mallat and Zhong (1992), using a gradient wavelet $\Psi = \nabla\Lambda$ where Λ is a smoothing kernel.

From an algorithmic point of view, when Λ is a Gaussian kernel, denoting N the size of the data (for an image $n \times n$, $N = n^2$), this formulation can be computed in $O(N)$. Practically, we compute at each considered scale s the gradient magnitude $Mf(.,.,s)$ and the gradient orientation $Af(.,.,s)$. Then multiscale MC are points $(x(s), y(s))$ where $Mf(.,.,s)$ attains a local maximum along the direction given by $Af(.,.,s)$.

Empirically, multiscale MC can be connected towards finer scales: in the Scale-Space domain, MC generate uninterrupted lines when the scale decreases (for details on wavelet maxima lines, see Mallat and Hwang 1992; Mallat 1998; Damerval and Meignen 2007). This leads to the notion of Canny's Maxima Line (CML), defined as

$$(x(s), y(s), s, Mf(x(s), y(s), s), Af(x(s), y(s), s)) \tag{15}$$

for all $s \in I$, where $I =]0, z[$, $z > 0$ or $z = +\infty$. Such CML appear as relevant since when f is α-Lipschitz, we hold along a CML (Mallat and Zhong 1992):

$$\log Mf(x(s), y(s), s) \leq \alpha \log s + C \tag{16}$$

Since this inequality is a quasi-equality at the finest scales, a precise estimation of α at the origin of a CML can be obtained by performing a linear regression over three fine scales, see Fig. 5 (see also Appendix B for the link between wavelets and regularity α).

For robustness purpose, we discard CML that do not cross sufficient scales and also apply a light thresholding to the response at a fine scale, which leads to a set of relevant CML. Finally the algorithm we use can be summarized as follows:

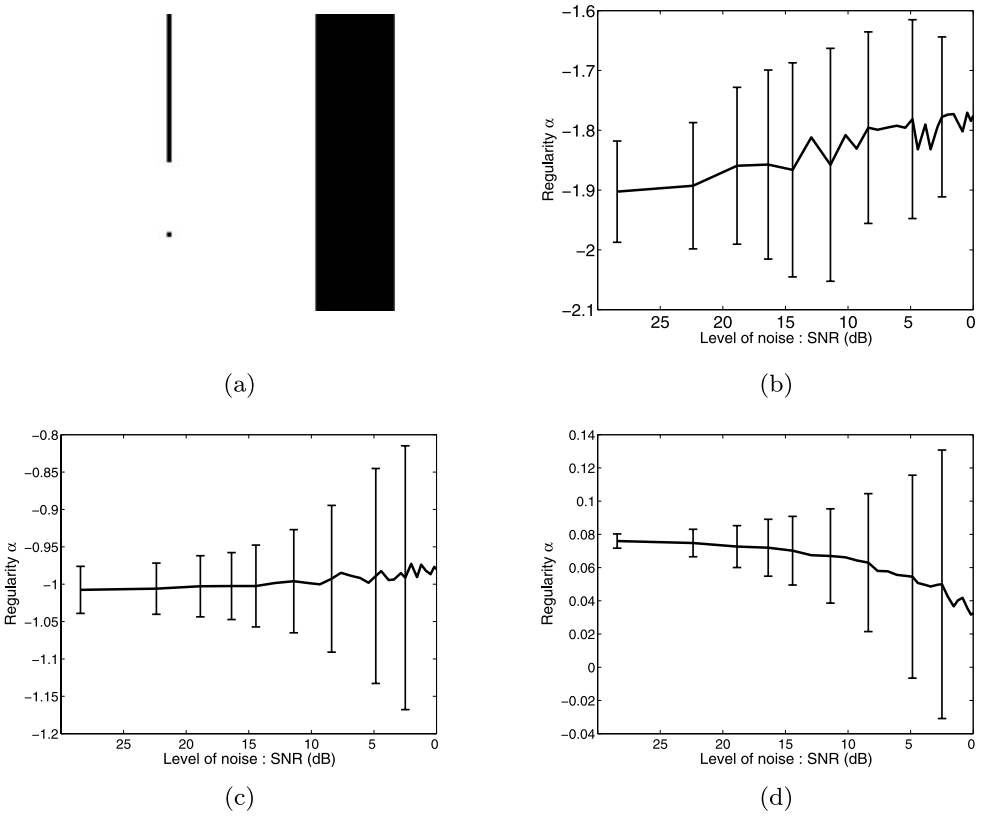

Fig. 6 Influence of noise on the regularity α. (**a**) Geometric image containing *an isolated point*, *a line* and *an edge*. (**b**)–(**d**) Evolution of α as the level of noise increases (in terms of mean and standard deviations) for: an isolated point (**b**), a line (**c**) and an edge (**d**)

1. Construct all CML associated to the image f
2. Select relevant CML
3. Along each relevant CML, estimate α

Denoting n_{CML} the number of relevant CML, this approach associates to the image f a set

$$\left\{(x_i, y_i, \alpha_i) \in \mathbb{R}^3, 1 \leq i \leq n_{CML}\right\} \tag{17}$$

where (x_i, y_i) corresponds to a contour point (origin of a CML) and α_i is the estimated regularity α at (x_i, y_i). We point out that the selection of relevant CML is not a very strict criterion. Hence we obtain values of regularity α at many contour points. Note also that since contour point correspond to singularities where f may be not differentiable, the obtained values can be negative: typically an edge leads to $\alpha = 0$, a line to $\alpha = -1$, and an isolated point to $\alpha = -2$.

In our numerical experiments, we use integer scales ($s = 1, 2, 3, \ldots$), select CML that cross at least 4 scales, and impose $Mf(., ., 1) > thresh$ ($thresh = 0.04$ for an intensity between 0 and 1). Finally, the estimation of α is based on the scales $s = 1, 2, 3$. Other choices are possible, but one should keep in mind that choosing too fine scales can lead to numerical instabilities of the linear regression and that choosing too coarse scales reduces drastically the number of CML (and thus the number of extracted features). For illustration purpose, we represent examples of such contour points on Fig. 4. Let us now study the robustness of this algorithm, first in controlled situations (geometric images) and then in more general cases (natural images).

4.2 Robustness to Noise and to Geometric Deformations

We consider here geometric images for which the regularity α is known *a priori*. In these controlled situations, we show that the estimation of α is robust: in the presence of noise, and also under geometric transformations such as rotation or scale change.

First, we consider a 128×128 geometric image containing an isolated point, a line and an edge (Fig. 6(a)), to which we add a white Gaussian noise (100 simulations for each level of noise, quantified by the signal-to-noise ratio SNR). Then we estimate the regularity α at these contour points, and, for each kind of contour (isolated point, line, edge), we represent the evolution of the regularity α with respect to the SNR—in terms of average and standard deviation (Figs. 6(b)–(d)). In the absence of noise, the computed values are exactly those expected, and as the level of noise increases, the quality of the estimation slightly decreases. More precisely, the isolated point is clearly affected by noise addition, whereas the estimation of α is more robust in the case of a line or an edge.

Secondly, we consider a 128×128 image representing a line (resp. an edge), rotated by θ ($0 \leq \theta \leq 90°$). Then we compute values of regularity α and, in terms of average and

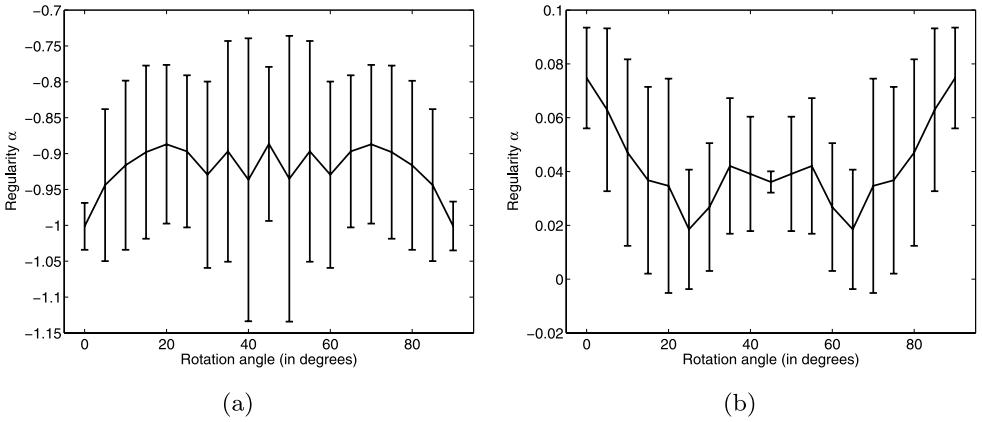

Fig. 7 Influence of a rotation on the regularity α. Evolution of α with respect to the angle of rotation applied to a geometric image 128×128 representing (**a**) a line ($\alpha = -1$), (**b**) an edge ($\alpha = 0$)

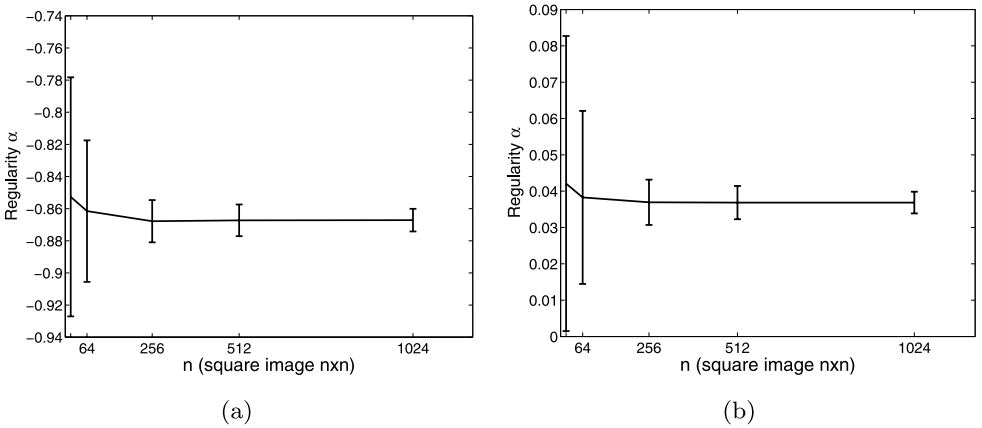

Fig. 8 Influence of a scale change on the regularity α. Evolution of α with respect to n on $n \times n$ geometric images representing either (**a**) a diagonal line ($\alpha = -1$) or (**b**) a diagonal edge ($\alpha = 0$)

standard deviation, we represent the evolution of the regularity α with respect to the angle θ (Figs. 7(a), (b)). We observe that a rotation has a limited influence on the estimation of α. More precisely we observe a certain bias, which is limited for the line (less than 0.11 on average) and small for the edge (less than 0.08 on average). So the regularity α appears robust to rotation.

Thirdly, we consider a $n \times n$ image representing a diagonal line (resp. a diagonal edge), and study the influence of the size n (corresponding to a scale change, and also to the study of sampling effect). Then we represent the evolution of α with respect to the size n—in terms of average and standard deviation (Figs. 8(a), (b)). Asymptotically, we observe a bias of 0.04 for $\alpha = 0$ (edge) and 0.13 for $\alpha = -1$ (line), and this bias is stable as n varies. (α being more precisely estimated for larger values of n.) Globally, the influence of a scale change on α is thus limited (in the case of a horizontal or vertical line, the estimated value of α corresponds always to the theoretical value).

Let us discuss briefly this bias. Since a rotated contour is not ideal (due to pixelization), it is difficult to evaluate accurately the theoretical value of α. Moreover, to compute the regularity α, we use the direction of the gradient as the direction of maximum irregularity. In certain cases, this direction may be inaccurately estimated, resulting in a larger value than the expected one, explaining this positive bias.

4.3 Detailed Study in the Case of an Affine Deformation

We study here the robustness of the estimation of α in the case of natural images, for which values of regularity α are computed at detected contour points. For that purpose we consider an original image X_0 (Fig. 9(a)) and a deformed image X_1 (Fig. 9(b)) related by a known affine deformation. Contours and associated values of α are computed independently for X_0 and X_1. Since we know the affine deformation, we can carry out point-to-point correspondences and thus to compare the values of regularity between both images.

Given X_0 and X_1, the preceding algorithm leads to two sets of contour points with associated values of regularity α:

$$\begin{cases} S_0 = (x_i^0, y_i^0, \alpha_i^0)_i \in I_0 & \text{(see Fig. 9(a'))} \\ S_1 = (x_j^1, y_j^1, \alpha_j^1)_j \in I_1 & \text{(see Fig. 9(b'))} \end{cases}$$

Afterwards, we project the points $(x_i^1, y_i^1)_i \in I_1$ into the coordinates of X_0, and we carry out correspondences be-

tween the sets S_1 and S_0. At this step, we have two possible choices:

– *Exact matches* (EM) for which a projected point of S_1 corresponds exactly to a point of S_0;

– *Approximate matches* (AM), by tolerating an error of 1 pixel; this gives rise to a larger number of correspondences and a loss of accuracy, as we will see further.

This allows to compare the computed values of α: given a correspondence between $(x_i^0, y_i^0, \alpha_i^0)$ and $(x_j^1, y_j^1, \alpha_j^1)$ (either EM or AM), we define the error on α as

$$d_\alpha = |\alpha_i^0 - \alpha_j^1| \tag{18}$$

Then, a match is said correct if $d_\alpha < \epsilon$ (where ϵ is a tolerated error on α). Now, let us compare S_0 and S_1.

First, we represent the density associated to the regularity α: on Fig. 10(a), considering all points of S_0; on Fig. 10(b), considering points of S_0 for which there is a correspondence (AM) with S_1. A statistical test shows the two densities are similar (according to a Wilcoxon test, with 10% significance level, Lehmann 1997). This shows that the considered geometric transformation does not affect a particular kind of contour point. Moreover we observe that 95% of the computed values are within $[-1.4, 0.8]$, for which the estimation of α was shown robust in the previous section.

Secondly, we represent on Fig. 11(a) the proportion of correct matches depending on ϵ, for both exact and approximate matches (we detect 57829 contour points in the original image and 55045 in the deformed image, which leads

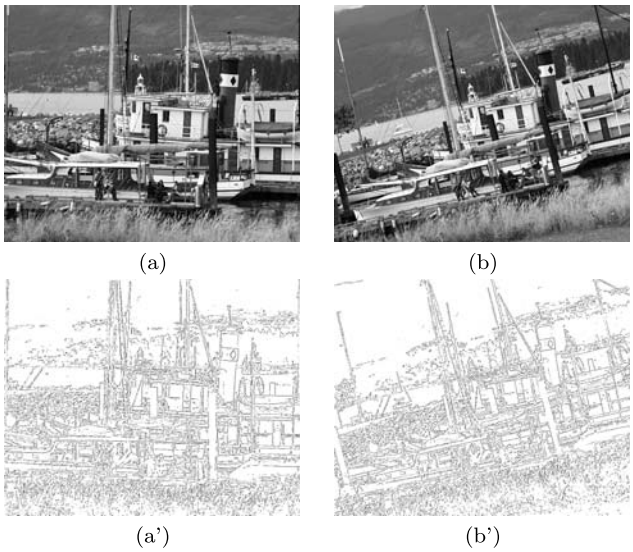

Fig. 9 (**a**), (**b**) Two images related by a known affine deformation; (**a'**), (**b'**) Detected contour points

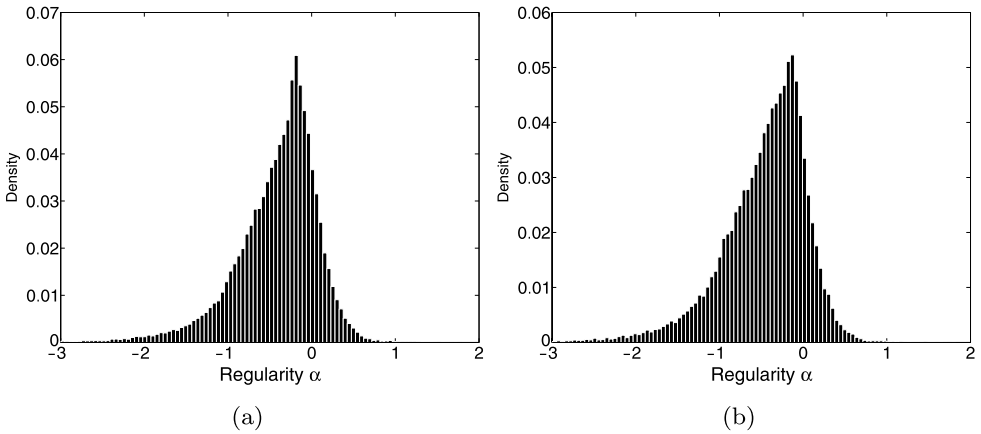

Fig. 10 Density of the regularity α associated to contour points of the image of Fig. 9(a), regarding: (**a**) all contour points; (**b**) only contour points associated to a correspondence

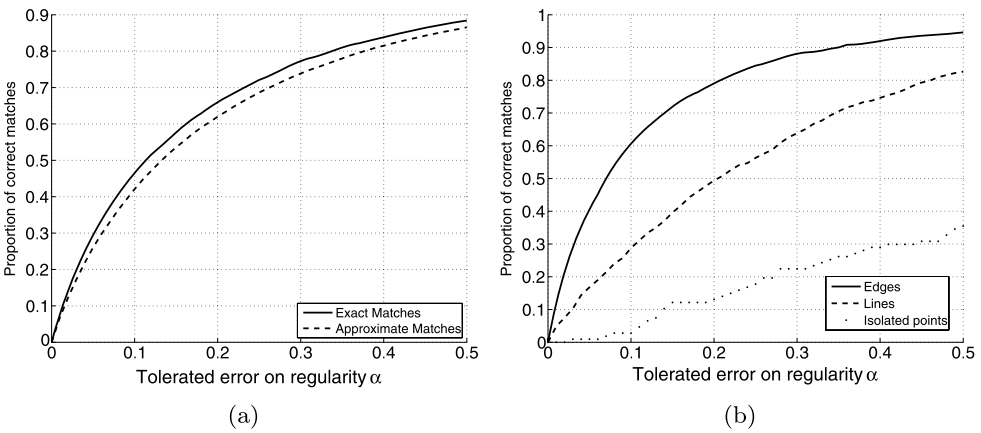

Fig. 11 (**a**) Errors on the estimation of α for exact matches (EM) and approximate matches (AM); (**b**) Errors on the estimation of α for exact matches, associated respectively to edges, lines, and isolated points

Table 1 How imaging conditions are performed and quantified

Imaging condition and sequence	Performed by	Quantified by	X_0	X_1	X_2	X_3	X_4	X_5
Scale change/rotation	camera zoom and orientation	scale factor						
ZoomRotation1, Fig. 12(a)			1	1.2	1.8	2.5	3	4
ZoomRotation2, Fig. 12(b)			1	1.1	1.4	1.9	2.4	2.8
Viewpoint change	camera position	angle of viewpoint						
Viewpoint1, Fig. 12(c)			0°	20°	30°	40°	50°	60°
Viewpoint2, Fig. 12(d)								
Blur	camera focus	level of blur						
Blur1, Fig. 12(e)		arbitrary values	1	2	3	4	5	6
Blur2, Fig. 12(f)		(from sharp to blurred)	1	2	3	4	5	6
JPEG compression	using a software	rate of compression						
Jpeg, Fig. 12(g)			0%	60%	80%	90%	95%	98%
Illumination change	camera aperture	level of light						
Light, Fig. 12(h)		arbitrary values (from light to dark)	1	2	3	4	5	6

to 38760 AM and 5924 EM). As the tolerated error on α increases, the proportion of matches identified as correct increases. Besides, we observe that the results are only slightly better for EM than AM, so that it can be relevant to consider AM (since their number is significantly larger).

Thirdly, we represent on Fig. 11(b) the proportion of correct matches depending on ϵ for exact matches associated with edges ($-0.3 < \alpha < 0.3$, 2466 points), lines ($-1.3 < \alpha < -0.7$, 1174 points) and isolated points ($-2.3 < \alpha < -1.7$, 107 points) respectively. This shows that the robustness of the regularity α is better for edges and lines than for isolated points. This is due to the fact that structures like edges and lines are more stable than isolated points under different viewing conditions (in this regard, note that isolated points are scarce compared to lines and edges). So globally the robustness of α increases with the value of α. Besides, most of detected points correspond to values of regularity α for which the estimation is robust (see Fig. 10).

Conclusions The proposed method allows to detect contour points and to compute associated values of regularity α. We studied the quality of the estimation of the regularity α, considering geometric images (noise addition, rotation, scale change) and also natural images (affine deformation). From these results we conclude that the regularity α estimated at contour points is a robust feature ensuring quantity (large number of contour points) and quality (robustness to an affine deformation). Let us now evaluate the robustness of this feature in a more general context, when various transformations are applied to natural images.

5 Study of More General Transformations

5.1 Data Set

We consider 8 sequences, each one consisting of 6 images $(X_k)_{0\leq k\leq 5}$. The sequences are called ZoomRotation1, ZoomRotation2, Viewpoint1, Viewpoint2, Blur1, Blur2, Jpeg and Light. For each sequence, the 6 images represent a given scene viewed under a certain imaging condition (see Table 1). For instance, considering the sequence Viewpoint1 (see Fig. 13), each image $X_k, 1 \leq k \leq 5$ corresponds to a change of viewpoint applied to the original image X_0. For illustration purposes, the images X_0 and X_5 associated to the different sequences are represented on Fig. 12 (see also Table 1).

The relevance of these sequences lies in different aspects. First they represent various objects: textured scenes (repeated textures, like Figs. 12(a), (d)) and structured ones (homogeneous regions with edges boundaries, like Figs. 12(b), (c)). Secondly the imaging conditions are wide-ranging: geometric deformations and specific transformations (like JPEG compression). Thirdly the level of these transformation can be significant (scale change up to 4, angle of viewpoint up to 60°, JPEG compression rate up to 98%).

These sequences were originally designed for the evaluation of interest point detectors and affine covariant features (Mikolajczyk et al. 2005); the complete data set we used is available at the following url: http://www.robots.ox.ac.uk/~vgg/research/affine.

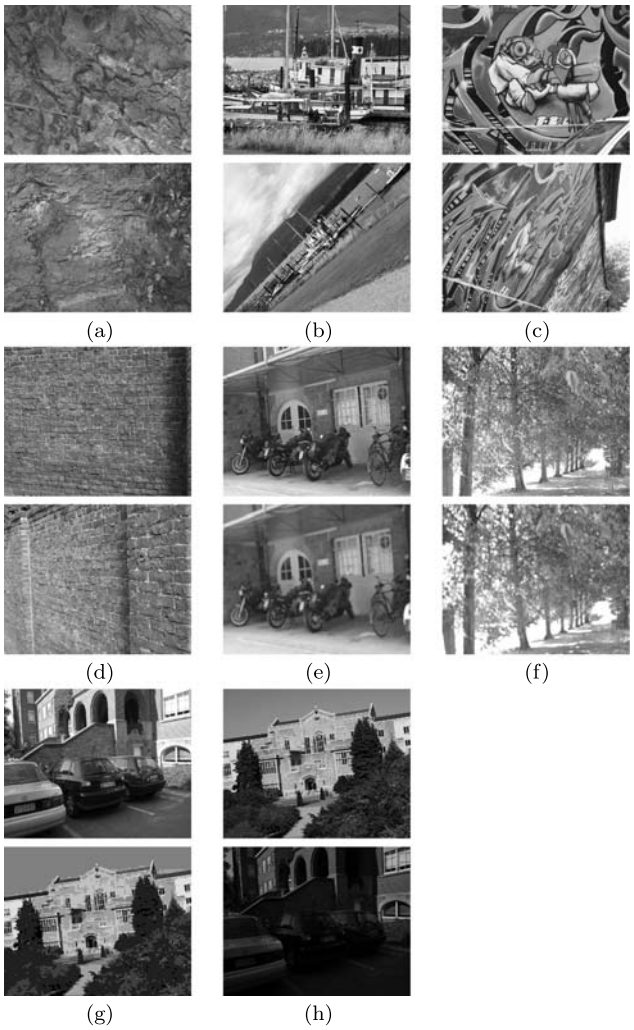

Fig. 12 Sample of data set, representing X_0 (reference image) and X_5 (highest level of transformation) associated to: (**a**), (**b**) Scale change and rotation; (**c**), (**d**) Viewpoint change; (**e**), (**f**) Blur; (**g**) JPEG compression; (**h**) Illumination change

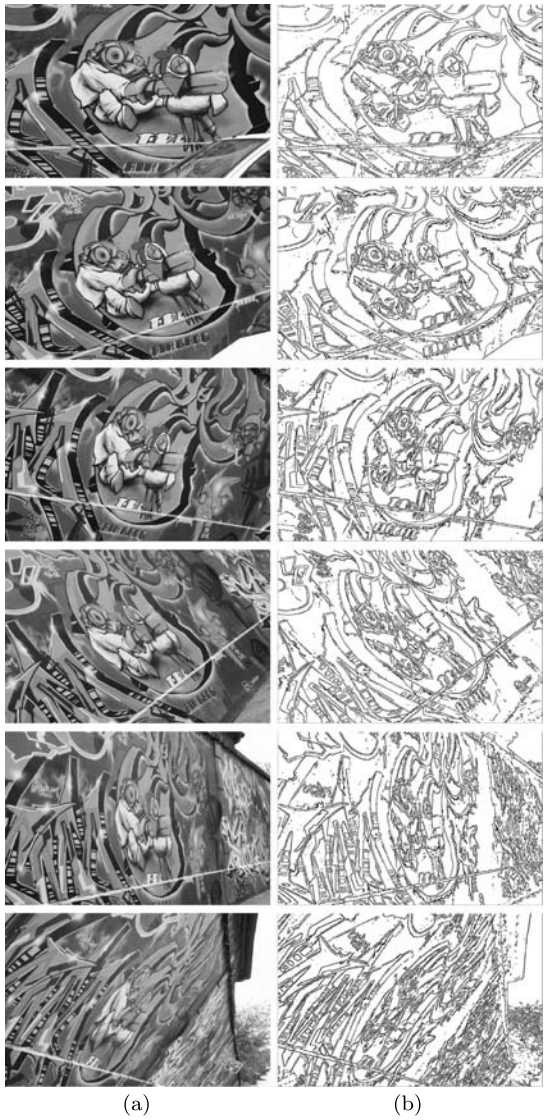

Fig. 13 (**a**) Complete sequence Viewpoint1 (images $X_0, \ldots, X_5$); (**b**) Associated contour points

5.2 Methodology

For a given set of images $(X_k)_{0\leq k\leq 5}$ associated to a sequence (viewpoint change for instance, see Fig. 13), we carry out the following procedure:

1. For each image $(X_k)_{0\leq k\leq 5}$, detect contour points and compute the associated values of regularity α: $p_i^k = (x_i^k, y_i^k, \alpha_i^k), 0 \leq i \leq n_k$
2. For fixed k $(1 \leq k \leq 5)$, determine a set of $\mathcal{C}^k$ of point-to-point correspondences between contour points of X_0 and X_k (thanks to the known homography between these images)

$$\mathcal{C}^k = \left\{ \begin{array}{l} (p_i^0, p_j^k) \text{ matched} \\ \text{according to a geometric criterion} \end{array} \right\} \tag{19}$$

This leads to a certain number of correspondences (NC) $\#\mathcal{C}^k$, and the repeatability score (RS):

$$R_k = \frac{\#\mathcal{C}^k}{\min(n_0, n_k)} \tag{20}$$

3. Select the subset $\mathcal{C}_\epsilon^k$ of correspondences for which regularities are sufficiently close (according to a parameter $\epsilon > 0$)

$$\mathcal{C}_\epsilon^k = \left\{ (p_i^0, p_j^k) \in \mathcal{C}^k, d_\alpha = |\alpha_i^0 - \alpha_j^k| < \epsilon \right\} \tag{21}$$

and compute the matching score:

$$S_k = \frac{\#\mathcal{C}_\epsilon^k}{\#\mathcal{C}^k} \tag{22}$$

representing the proportion of correct matches.

Remarks In step 2, the matches based on a geometric criterion can be either exact or approximate (as described in Sect. 4.3); we study both. In step 3, one has to choose ϵ, which is a parameter representing a tolerated error on α. This value plays a role similar to the error of a statistical hypothesis test (Lehmann 1997), so there is a trade-off between too low and too high values. We use here $\epsilon = 0.3$, for which almost 80% of the matches are deemed correct in the case of the affine deformation studied in Sect. 4.3 (see Fig. 11).

Now, considering the 8 sequences, for both approximate and exact matches, we represent the number of correspondences $\#\mathcal{C}^k$ with respect to k (Fig. 14) and the repeatability score R_k with respect to k (Fig. 15). These quantities assess the robustness of the contour points detector. Besides, we display on Fig. 16 the curves (k, S_k). This allows to assess the robustness of the estimation of the regularity α based on the CML, which is the main objective of this paper.

Note on the analysis of the results: a detector is efficient if it obtains high scores, and if these scores are stable. Here stability means that the score does not fall down as the level of the transformation increases.

5.3 Performance of the Contour Point Detector

The performance of the contour point detector can be evaluated in absolute terms—the number of correspondences (NC)—and in relative terms—the repeatability score (RS). Figure 14 shows that for all sequences, the NC globally decreases as the level of transformation increases. We observe that this decay is less pronounced for EM; so these are more robust, but in limited number compared to AM. Note also that the NC remains significant for all transformations and all associated levels. Regarding repeatability (Fig. 15), the RS concerning EM are low for all sequences. This is mainly due to the fact that the condition of exact correspondence is very strict. However, concerning AM, the RS become significantly higher. This shows the relevance of the tolerance of one pixel for point-to-point exact correspondence. Let us consider the results on AM more in detail, describing the level of the RS and its stability as the level of transformation increases.

Scale change and rotation: the RS remains low for textured scene (Fig. 15(a)). Note that this scene is very unstructured, and that contour points embedded in textures are less stable than those associated with sharp edges. The RS is higher for structured scene, for which it decreases steadily (Fig. 15(b)).

Viewpoint change: the RS decreases steadily, being higher for textured scene (Figs. 15(c), (d)). Note that this textured scene (Fig. 15(d)) contains some sharp edges, compared to Fig. 15(a). Globally the RS is higher than for scale change and rotation, mainly because these scenes contain sharper edges and lines.

Blur: the RS decreases rapidly for the structured scene (Fig. 15(e)) while we obtain good and stable RS for the textured scene (Fig. 15(f)).

JPEG compression (Fig. 15(g)): high RS, decreasing only for a high compression rate. In this regard, note that image compression methods were recently designed so as to preserve contours (Mallat 1989; Candes and Donoho 2000; Mallat and Peyré 2007).

Illumination change (Fig. 15(h)): high and stable RS. Contours are almost unaffected by a change in illumination.

General conclusions: the detector is efficient for intermediate levels of scale change/rotation, viewpoint change, and for high levels of JPEG compression and illumination change. Besides, images with sharp edges tend to yield better results. We emphasize that the RS presented here are globally satisfying since we use strict correspondence rules (maximum tolerance of one pixel). Besides, scenes with sharp edges tend to yield better results. We emphasize that the RS presented here are globally satisfying, since we use strict correspondence rules (maximum tolerance of one pixel). In addition, we point out that geometric deformations can lead to major occlusions of certain parts of the image, which is not the case of the other transformations (like blur, JPEG compression and illumination change). This explains why RS are lower on Figs. 15(a)–(c) than on Figs. 15(e)–(h). Eventually, we observe that the numbers of correspondences is often large for EM and very large for AM (Fig. 14). In conclusion, the algorithm based on CML detects many contour points, and among them, a significant number is robust to various transformations of the image.

Note on Figs. 14–16: each graph of these figures describes the performance of our method for one particular sequence associated to a certain image transformation (see Fig. 12 and Table 1). For instance, Fig. 14(c), Fig. 15(c) and Fig. 16(c) refer to the sequence Viewpoint1 (associated to a viewpoint change), represented partially on Fig. 12(c) (and comprehensively on Fig. 13). For all graphs, each point corresponds to a comparison between two images.

5.4 Evaluation of the Robustness of Regularity α

On the basis of the matching scores (see Sect. 5.2), we can evaluate how robust is the estimation of α under various imaging conditions (Fig. 16). For clarity purpose, let us recall that once CML are constructed, the origin of a CML gives a contour point where regularity α is estimated. The matching score gives the proportion of correspondences (between two images) for which computed values are close. Globally, we observe that the score tends to decrease as the level of the transformation increases. Note also the scores for EM and AM are close, and that EM yield better results than AM (as we pointed out in Sect. 4.3, see Fig. 11).

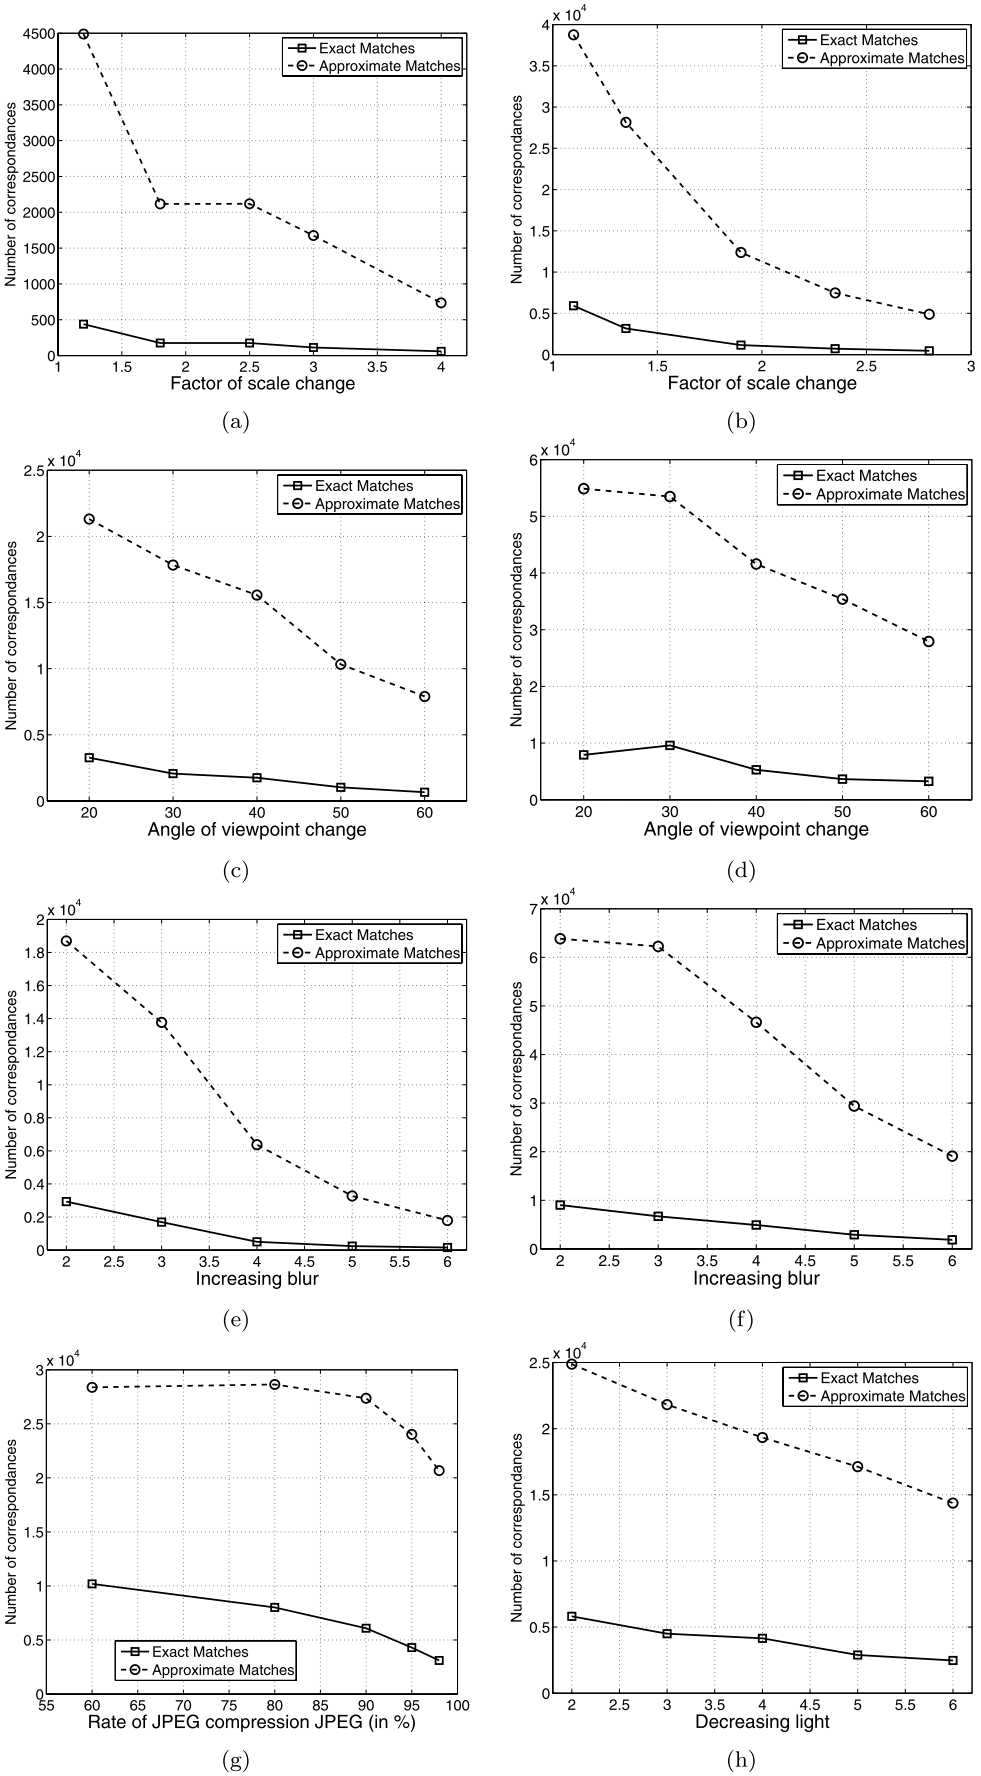

Fig. 14 (**a**)–(**h**) Number of correspondences (NC) of the contour point detector, associated to the sequences of Fig. 12(a)–(h)

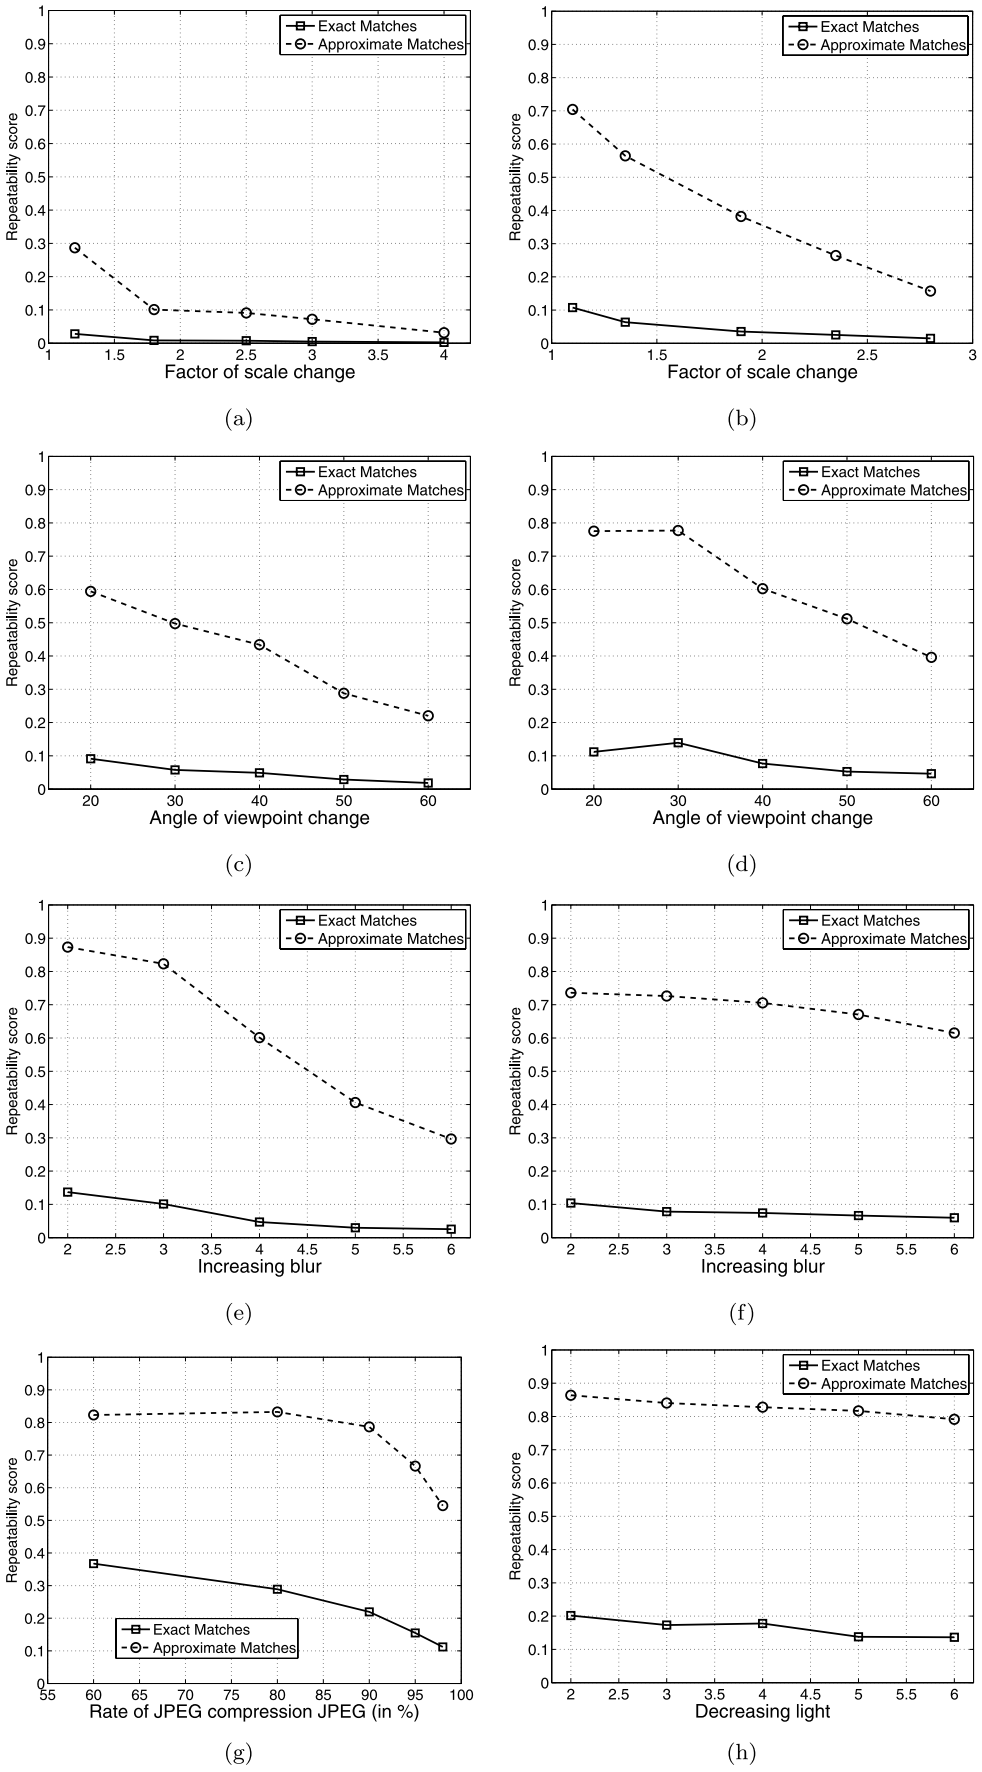

Fig. 15 (**a**)–(**h**) Repeatability scores (RS) of the contour point detector, associated to the sequences of Fig. 12(a)–(h)

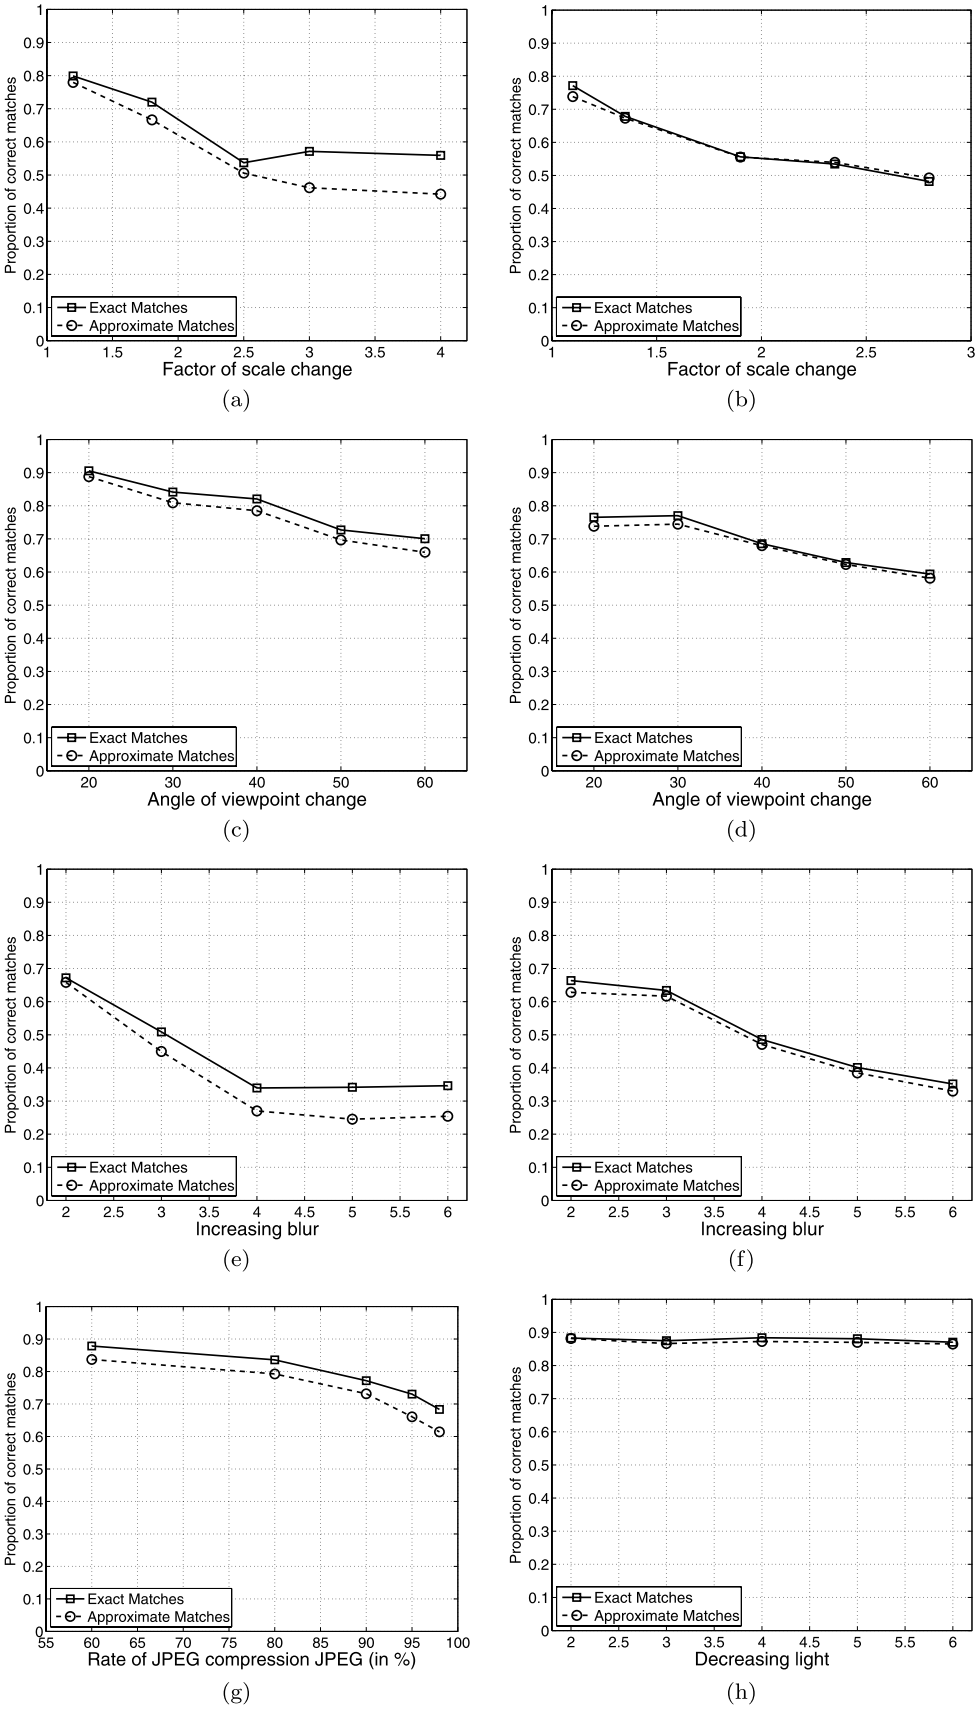

Fig. 16 (**a**)–(**h**) Matching scores associated to the different sequences of Fig. 12(a)–(h). In our experiments, the tolerance parameter on α is set to $\epsilon = 0.3$

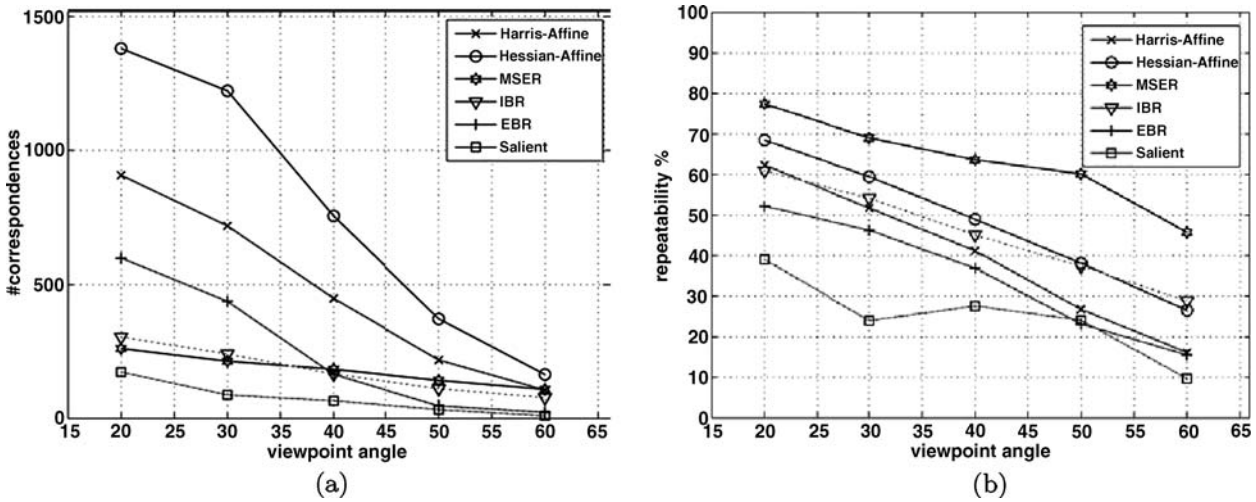

Fig. 17 Robustness of state-of-the-art detectors for the sequence Viewpoint1, (each curve corresponds to a detector): (**a**) Number of correspondences; (**b**) Repeatability scores

We do not observe a significant difference between textured (Figs. 16(a), (d)) and structured scenes (Figs. 16(b), (c)); however, the structured scene of Fig. 16(c) leads to better results, due to the presence of sharp edges and lines. Now, let us comment the results more in detail.

Scale change and rotation: for both textured and structured scenes (Figs. 16(a), (b)), the performance decreases overall (from 0.8 to 0.5). Since sequences ZoomRotation1 and ZoomRotation2 correspond to significant scale changes and rotations, the obtained results are satisfactory.

Viewpoint change: the performance decreases moderately, remaining high for the structured scene (between 0.9 and 0.7, Fig. 16(c)) and good for the textured scene (between 0.8 and 0.6, Fig. 16(d)). In this regard, note that both sequences Viewpoint1 and Viewpoint2 contain distinct contour, thus moderately affected by a viewpoint change.

Blur: the performance decreases rapidly (from 0.7 to 0.3) for both structured scenes and textured ones (Figs. 16(e), (f)). This is not surprising to obtain these average results since the blur modifies the regularity α, as we pointed out in Sect. 3.2 (edges are smoothed).

JPEG compression (Fig. 16(g)): the performance remains high (between 0.9 and 0.6), decreasing steadily. So JPEG artifacts have little impact on regularity α, even if JPEG compression tend to blur sharp edges.

Illumination change (Fig. 16(g)): very high performance (stable, close to 0.9). This is related to the invariance of the regularity α to an illumination change (see Sect. 3.2).

General conclusions: the regularity α appears as very robust to light change and JPEG compression, and less to image blur. Concerning geometric deformations, we obtain very good results for viewpoint change (especially structured scenes) and satisfying ones for scale change and rotation. In addition, we observe that for all transformations except blur, the performance does not fall down, even for a high level of transformation. This emphasizes the fact that the regularity α is characteristic of the kind of contour. Eventually, note that there is a balance between quantity—larger number of AM than EM, better repeatability for AM—and quality—better estimation of α for EM than AM. More precisely this balance is in favor of AM since scores are only slightly inferior for AM than for EM.

5.5 Comparison with Interest Point Detectors

We compare here certain aspects of our detector to the best interest point detectors (Mikolajczyk et al. 2005). We represent on Fig. 17 some curves which evaluate the performance of such methods. More precisely these represent the number of correspondences and the repeatability scores of different detectors, for one given sequence (here it is Viewpoint1, represented on Fig. 13(a)). The curves concerning all sequences studied in this paper can be found in Mikolajczyk et al. (2005).

First note that the extracted features are different: local features (regions) for these methods, and pointwise features for our approach. We focus on position and pointwise regularity α, whereas those are based on position and characteristic scale. While a contour point is characterized by its regularity α (both theoretically and practically), the content of a region is characterized by a descriptor (only practically). In terms of number of correspondences we obtain better results, due to more points detected (typically 10000 points vs. 1000 regions). In terms of repeatability, we obtain worse results for EM, which is not surprising: tolerated errors between interest regions can be high (overlap errors up to 40%) whereas the condition of EM is very strict. However, the results obtained for AM are globally similar to those of the best existing detectors. Besides, we can also compare matching scores of these detectors (criterion based

on a descriptor) to those presented here (criterion based on a tolerated error on α). Since the two criteria are different, we can only draw conclusions from the shape of the curves. In this regard, we observe our method is more stable than interest point detectors: as the level of transformation increases, the performance declines slower overall. More precisely, our method appears: more stable for JPEG, Viewpoint1, Viewpoint2 and ZoomRotation2; less stable ZoomRotation1 and Blur1; equivalent for Light, Blur2.

5.6 Discussion

Let us discuss now some choices we made in our approach. First, there exist methods based on wavelet transforms that allow to estimate the regularity α at any point of the image. Here we focused on contour points, since they ensure good localization and appear as robust to various transformations of the image (as discussed in Sect. 3.2). Besides, the repeatability performance of the contour point detector can be improved: in the algorithm presented in Sect. 4, one can impose strict conditions on the selected CML (by increasing the number of crossed scales, and/or the value of the threshold *thresh*); this leads to a limited and more robust set of points. Here we used a light thresholding so as to obtain a large number of estimated values of the regularity α. Finally, the matching performance can also be improved by applying a threshold on α (selecting only edges for instance). Here, we considered all values of the regularity α, showing that a significant proportion of the computed values of α correspond to values for which the estimation is robust (as we pointed out in Sect. 4.3, see also Figs. 11(a)–(b)).

6 Conclusions

In this paper the Lipschitz regularity was studied from different aspects. From a theoretical point of view, we showed results on the influence of certain transformations on α-Lipschitz 1D functions. In a 2D context, we explained why certain transformations of the image do not change the pointwise regularity α. From a practical point of view, we recalled how the regularity α can be numerically estimated at contour points, thus giving a characteristic value of the kind of contour.

The main contribution of our work lies in quantifying the robustness of this estimated value of α. For that purpose we proposed an evaluation procedure which allows to compare the values of the regularity α between two images related by a known homography. While the number of correspondences and the repeatability scores quantify the performance of the contour point detector, the matching scores reflect the robustness of the regularity α. Concerning the contour point detector, the obtained repeatability scores appear as satisfactory, since the criterion regarding point-to-point correspondences is very strict. As a matter of fact, the number of correspondences remains large for any level of the considered transformations, which is desirable in the applications.

The regularity α appears as a robust feature, especially in case of geometric deformations (such as viewpoint change, scale change and rotation) and of JPEG compression and illumination change. We compared our approach to another classical method for feature extraction—the detection of interest regions—pointing out their differences and showing that the robustness of the regularity α was similar to the best state-of-the-art methods. We underline the originality of the regularity α: whereas classical methods focus on local features that co-vary under certain transformations, our approach focus on point-wise features that are directly invariant to various transformations (geometric and photometric).

7 Perspectives

From a practical point of view the regularity α stands out as interesting for the applications: one the one hand it gives a characteristic value, robust and repeatable under various imaging conditions; on the other hand, the approach presented here allows to extract many features and also yields a fast computation of the regularity α. Let us point out the perspectives of this work. First, a new interest region detector may be defined, for instance by clustering contour points on a criterion based on both regularity α and geometry (contour points tend to have connectivity relations). Secondly, the regularity α could be used complementary to interest region descriptors: the estimated regularity α at all contour points within a given region may help to characterize the content of the region (this comes from the fact the presented algorithm can evidence a great number of pointwise features). Eventually, the regularity α appears as an interesting additional feature, complementary to other existing local features: an integration of the regularity α in existing methods will certainly improve their performance.

From a theoretical point of view, let us point out the generalization of measures like the regularity α. For instance, in the field of junction detection, the regularity α can only measure the kind of edge that meet at a junction. So it would be interesting to measure the arity of a junction (how many edges meet) and additionally get the actual geometric configuration.

Software To compute the different scores presented here, an evaluation procedure (Matlab code) is available at the url: http://www2.cs.kuleuven.be/~christod/Software.

Acknowledgements We would like to thank the referees for useful comments and suggestions.

Appendix A: Proofs of Propositions 1–4

Here we detail the proofs of Propositions 1–4, quoted in Sect. 2 (study in the monodimensional case).

Proposition 1 *Let $x_0 \in \mathbb{R}$ and $f, g : \mathbb{R} \to \mathbb{R}$ related by $g(x) = f(u^{-1}(x))$ where $u : \mathbb{R} \to \mathbb{R}$. We assume that u is C^∞, invertible and linear (so that u^{-1} is linear)*

$$\forall \alpha > 0, \quad f \ \alpha\text{-Lipschitz at } x_0 \implies g \ \alpha\text{-Lipschitz at } x_0 \tag{3}$$

Proof of Proposition 1 Let us distinguish the cases $\alpha \in]0, 1[$ and $\alpha > 1$. When $\alpha \in]0, 1[$, there exist $A > 0$ and $K > 0$ such that

$$\begin{aligned} |g(x) - g(x_0)| &= |f(u^{-1}(x)) - f(u^{-1}(x_0))| \\ &\leq A|u^{-1}(x)) - u^{-1}(x_0)|^\alpha \\ &\leq K|x - x_0|^\alpha \end{aligned}$$

So assertion (3) is valid when $\alpha \in]0, 1[$.

When $\alpha > 1$ (non integer), first note there exists a polynomial P_n of order n satisfying

$$\begin{aligned} |g(x) - P_n(u^{-1}(x))| &= |f(u^{-1}(x)) - P_n(u^{-1}(x))| \\ &\leq A|u^{-1}(x) - u^{-1}(x_0)|^\alpha \\ &\leq A_1|x - x_0|^\alpha \end{aligned}$$

Moreover, as $P_n(u^{-1}(\cdot))$ is C^∞, it is also α-Lipschitz. So there exists a polynomial Q_n of order n and $A_2 > 0$ so that

$$|P_n(u^{-1}(x)) - Q_n(x)| \leq A_2|x - x_0|^\alpha$$

Thus there exists $K > 0$ so that

$$\begin{aligned} |g(x) - Q_n(x)| &\leq |g(x) - P_n(u^{-1}(x))| \\ &\quad + |P_n(u^{-1}(x)) - Q_n(x)| \\ &\leq A_1|x - x_0|^\alpha + A_2|x - x_0|^\alpha \\ |g(x) - Q_n(x)| &\leq K|x - x_0|^\alpha \end{aligned}$$

So assertion (3) is valid when $\alpha > 1$ (for $\alpha = n \in \mathbb{N}^*$, it is straightforward according to general theorems on C^n-functions). □

Proposition 2 *Let $x_0 \in \mathbb{R}$ and $f, g : \mathbb{R} \to \mathbb{R}$ related by $g(x) = c(x)f(x) + d(x)$ where $c, d : \mathbb{R} \to \mathbb{R}$. In the case $\alpha \in]0, 1[$ and c, d are C^1, or in the case $\alpha > 1$ and c, d are C^∞:*

$$\forall \alpha > 0, \quad f \ \alpha\text{-Lipschitz at } x_0 \implies g \ \alpha\text{-Lipschitz at } x_0 \tag{4}$$

Proof of Proposition 2 When $\alpha \in]0, 1[$, we can write

$$\begin{aligned} &|g(x) - g(x_0)| \\ &\quad = |c(x)f(x) - c(x_0)f(x_0) + d(x) - d(x_0)| \\ &\quad \leq |c(x)f(x) - c(x)f(x_0) + c(x)f(x_0) - c(x_0)f(x_0)| \\ &\qquad + |d(x) - d(x_0)| \\ &\quad \leq |c(x)| \cdot |f(x) - f(x_0)| \\ &\qquad + |f(x_0)| \cdot |c(x) - c(x_0)| \\ &\qquad + |d(x) - d(x_0)| \end{aligned}$$

As c is bounded on any compact interval I containing x_0, there exists C_1 such that for all $x \in I$:

$$\begin{aligned} |g(x) - g(x_0)| \leq C_1|f(x) - f(x_0)| \\ + |f(x_0)| \cdot |c(x) - c(x_0)| \\ + |d(x) - d(x_0)| \end{aligned}$$

Since c, d are continuous, they are also α-Lipschitz ($\alpha \in]0, 1[$), thus leading to

$$|g(x) - g(x_0)| \leq C|x - x_0|^\alpha \tag{23}$$

So assertion (4) is true when $\alpha \in]0, 1[$.

When $\alpha > 1$ (non integer), we can write

$$\begin{aligned} |g(x) - c(x)P_n(x) + d(x)| &= |c(x)| \cdot |f(x) - P_n(x)| \\ &\leq A_1|x - x_0|^\alpha \end{aligned}$$

Since $c(x)P_n(x) + d(x)$ is C^∞, it is also α-Lipschitz, there exists a polynomial $Q_n(x)$ of order n and A_2 such that

$$|c(x)P_n(x) + d(x) - Q_n(x)| \leq A_2|x - x_0|^\alpha$$

Thus there exists $K > 0$ so that

$$\begin{aligned} |g(x) - Q_n(x)| &\leq |g(x) - c(x)P_n(x) + d(x)| \\ &\quad + |P_n(u^{-1}(x)) - Q_n(x)| \\ &\leq A_1|x - x_0|^\alpha + A_2|x - x_0|^\alpha \\ &\leq K|x - x_0|^\alpha \end{aligned}$$

So assertion (4) is true when $\alpha > 1$ (for $\alpha = n \in \mathbb{N}^*$, it is straightforward according to general theorems on C^n-functions). □

Proposition 3 *Let T a distribution of finite order and T^u the associated distribution according to Definition 3. For all $\alpha \leq 0$,*

$$T \ \alpha\text{-Lipschitz} \implies T^u \ \alpha\text{-Lipschitz} \tag{7}$$

Proof of Proposition 3 There exists $f : \mathbb{R} \to \mathbb{R}$ so that

$$\forall \varphi \in \mathcal{S}(\mathbb{R}) \quad \langle T, \varphi \rangle = \langle f^{(n)}, \varphi \rangle$$

Thus we have

$$\langle T^u, \varphi \rangle = \langle T, (\varphi \circ u)u' \rangle = \langle f^{(n)}, (\varphi \circ u)u' \rangle \tag{24}$$

By putting $\psi = \varphi \circ u$, we can write

$$\begin{aligned}\langle T^u, \varphi\rangle &= \langle f^{(n)}, \psi u'\rangle = \langle f, (-1)^n [\psi\ u']^{(n)}\rangle \\ &= \left\langle f, (-1)^n \sum_{k=0}^{n} C_n^k \psi^{(k)} u^{(n-k+1)} \right\rangle \\ &= (-1)^n \sum_{k=0}^{n} C_n^k \langle f, \psi^{(k)} u^{(n-k+1)}\rangle \\ &= (-1)^n \sum_{k=0}^{n} C_n^k \langle f u^{(n-k+1)}, \psi^{(k)}\rangle \\ &= (-1)^n \sum_{k=0}^{n} C_n^k \langle f u^{(n-k+1)}, \psi^{(k)}\rangle \\ &= < (-1)^n \sum_{k=0}^{n} C_n^k \langle \left[f u^{(n-k+1)}\right]^{(k)}, \psi\rangle\end{aligned}$$

which leads to

$$\langle T^u, \varphi\rangle = \langle g, \psi\rangle \tag{25}$$

where g is defined as

$$\begin{aligned}g &= (-1)^n \sum_{k=0}^{n} C_n^k [f u^{(n-k+1)}]^{(k)} \\ &= (-1)^n \sum_{k=0}^{n} C_n^k \sum_{p=0}^{k} C_k^p f^{(p)} u^{(n-k+1+n-p)} \\ &= (-1)^n \sum_{k=0}^{n} C_n^k \sum_{p=0}^{k} C_k^p f^{(p)} u^{(n-k+1+n-p)}\end{aligned}$$

$$g = h u' \tag{26}$$

with h defined as

$$h = \sum_{k=0}^{n} \sum_{p=0}^{k} (-1)^n C_n^k C_k^p f^{(p)} u^{(2n-k-p)} \tag{27}$$

Restarting from (25), we obtain

$$\begin{aligned}\langle T^u, \varphi\rangle &= \langle g, \psi\rangle \\ &= \langle h u', \varphi \circ u\rangle \\ &= \langle h, (\varphi \circ u) u'\rangle \\ \langle T^u, \varphi\rangle &= \langle h \circ u^{-1}, \varphi\rangle\end{aligned} \tag{28}$$

Finally, note that in (27) the order of the derivatives of f is at most n. This is also true for $h \circ u^{-1}$, hence the primitive of order n of $h \circ u^{-1}$ has the same regularity as f (which is $(\alpha + n)$-Lipschitz). This shows T^u is α-Lipschitz. □

Proposition 4 *Let T a distribution of finite order, and T_1 the distribution given by $T_1 = cT + d$, assuming $c, d : \mathbb{R} \to \mathbb{R}$ are bounded and in $\mathcal{S}(\mathbb{R})$ (so that T_1 is well defined on $\mathcal{S}(\mathbb{R})$). For all $\alpha \leq 0$,*

$$T\ \alpha\text{-Lipschitz} \implies T_1\ \alpha\text{-Lipschitz} \tag{8}$$

Proof of Proposition 4 According to the properties of distributions:

$$\begin{aligned}\langle T_1, \varphi\rangle &= \langle cT + d, \varphi\rangle \\ &= \langle cT, \varphi\rangle + \langle d, \varphi\rangle \\ &= \langle T, c\varphi\rangle + \langle d, \varphi\rangle \\ &= \langle f^{(n)}, c\varphi\rangle + \langle d, \varphi\rangle \\ \langle T_1, \varphi\rangle &= \langle c f^{(n)} + d, \varphi\rangle\end{aligned}$$

Since the order of the derivatives of f in the expression $cf^{(n)} + d$ is at most n, T_1 has the same Lipschitz regularity as T. So T_1 is α-Lipschitz. □

Remark These invariance properties can be extended to the bidimensional case using distributions on $\mathbb{R}^2$. For instance, a line can be seen as the derivative in the sense of distributions of an edge (regularity $\alpha = 0$), therefore having regularity $\alpha = -1$.

Appendix B: Link Between Lipschitz Regularity and Wavelets. Application to the Computation of the Regularity α

Here we show the link between Lipschitz regularity and wavelets, by recalling a theorem shown by Jaffard and Meyer (1996), Mallat (1998) and the application to the computation of the regularity α. These 1D results can be generalized in 2D (and also in n-D, with n positive integer).

Proposition 5 *Let $f : \mathbb{R} \to \mathbb{R}$, $f \in L^2(\mathbb{R})$, and Wf its wavelet transform defined as: $\forall s > 0, \forall u \in \mathbb{R}$,*

$$Wf(u,s) = \frac{1}{s} \int_{\mathbb{R}} f(x) \psi\left(\frac{x-u}{s}\right) dx \tag{29}$$

(where ψ is a wavelet function).

If f is α-Lipschitz (with $\alpha \leq n$) at x_0 then there exists $A > 0$ such that: $\forall u \in \mathbb{R}, \forall s > 0$,

$$|Wf(u,s)| \leq A s^{\alpha} \left(1 + \left|\frac{u - x_0}{s}\right|^{\alpha}\right) \tag{30}$$

Conversely, if $\alpha < n$ is not an integer and there exist $A > 0$ and $\alpha' < \alpha$ such that: $\forall u \in \mathbb{R}, \forall s > 0$,

$$|Wf(u,s)| \leq A s^{\alpha} \left(1 + \left|\frac{u - x_0}{s}\right|^{\alpha'}\right) \tag{31}$$

then f is α-Lipschitz at x_0.

Application: computation of the regularity α Starting from (30) we get

$$|Wf(u,s)| \leq As^{\alpha} \tag{32}$$

so that

$$\log|Wf(u,s)| \leq \log A + \alpha \log s \tag{33}$$

and at fine scales this inequality becomes a quasi-equality (Mallat and Hwang 1992). So the regularity α can be estimated by performing a regression based on the formula

$$\log|Wf(u,s)| = \alpha \log s + C \tag{34}$$

More precisely α and C are computed by a regression of $\log|Wf(u,s)|$ over $\log s$ at fine scales (3 scales are sufficient for an accurate estimation).

References

Arneodo, A., Bacry, E., & Muzy, J. F. (1993). Singularity spectrum of fractal signals: exact results. *Journal of Statistical Physics*, *70*, 635–674.

Arneodo, A., Bacry, E., Jaffard, S., & Muzy, J. F. (1997). Oscillating singularities on cantor sets: a grand-canonical multifractal formalism. *Journal of Statistical Physics*, *87*, 179–209.

Arneodo, A., Bacry, E., Jaffard, S., & Muzy, J. F. (1998). Singularity spectrum of multifractal functions involving oscillating singularities. *Journal of Fourier Analysis and Applications*, *4*(2), 159–174.

Arneodo, A., Bacry, E., Jaffard, S., & Muzy, J. F. (1999). Oscillating singularities and fractal functions. Spline functions and the theory of wavelets. *CRM Proceedings and Lecture Notes*, *18*, 315–329.

Baumberg, A. (2000). Reliable feature matching across widely separated views. In *Computer vision and pattern recognition* (pp. 774–781).

Benassi, A., Cohen, S., Istas, J., & Jaffard, S. (1998). Identification of filtered white noises. *Stochastic Processes and Their Applications*, *75*(1), 31–49.

Bigot, J. (2003). Automatic landmark registration of 1d curves. In *Recent advances and trends in nonparametric statistics* (pp. 479–496). Amsterdam: Elsevier.

Brown, M., & Lowe, D. (2007). Automatic panoramic image stitching using invariant features. *International Journal of Computer Vision*, *14*(1), 59–73.

Candes, E., & Donoho, D. L. (2000). New tight frames of curvelets and optimal representations of objects with piecewise-c2 singularities. *Communications on Pure and Applied Mathematics*, *57*, 219–266.

Canny, J. (1986). A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *8*(6), 679–698.

Damerval, C., & Meignen, S. (2007). Blob detection with wavelet maxima lines. *IEEE Signal Processing Letters*, *14*(1), 39–42.

Deguy, S., Debain, C., & Benassi, A. (2000). Classification of texture images using multi-scale statistical estimators of fractal parameters. In *British machine vision conference*.

Florack, L., ter Haar Romeny, B., Koenderink, J., & Viergever, M. (1993). Cartesian differential invariants in scale-space. *Journal of Mathematical Imaging and Vision*, *3*(4), 327–348.

Iijima, T. (1962). Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bulletin of the Electrotechnical Laboratory*, *26*, 368–388.

Jaffard, S., & Meyer, Y. (1996). *Wavelet methods for pointwise regularity and local oscillations of functions*. Providence: American Mathematical Society.

Kadir, T., Zisserman, A., & Brady, M. (2004). An affine invariant salient region detector. In *European conference on computer vision* (pp. 404–416).

Kaplan, L. M., & Kuo, C. C. (1995a). Fast image synthesis using an extended self-similar model. *IEEE ICASSP-95*, *4*, 575–578.

Kaplan, L. M., & Kuo, C. C. (1995b). Texture roughness analysis and synthesis via extended self-similar model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *17*(11), 1043–1056.

Ko, H., & Lee, J. (2008). Gradient-based local affine invariant feature extraction for mobile robot localization in indoor environments. *Pattern Recognition Letters*, *29*, 1934–1940.

Koenderink, J. J., & Doorn, A. V. (1987). Representation of local geometry in the visual system. *Biological Cybernetics*, *55*, 367–375.

Lehmann, E. L. (1997). *Testing statistical hypotheses*. Berlin: Springer.

Lindeberg, T. (1994). *Scale space theory in computer vision*. Dordrecht: Kluwer Academic.

Lindeberg, T. (1998). Feature detection with automatic scale selection. *International Journal of Computer Vision*, *30*(2), 77–116.

Lindeberg, T., & Garding, J. (1997). Shape-adapted smoothing in estimation of 3-d shape cues from affine deformations of local 2-d brightness structure. *Image and Vision Computing*, *15*(6), 415–434.

Lowe, D. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, *60*(2), 91–110.

Mallat, S. (1989). A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *11*, 674–693.

Mallat, S. (1998). *A wavelet tour of signal processing*. San Diego: Academic Press.

Mallat, S., & Hwang, W. L. (1992). Singularity detection and processing with wavelets. *IEEE Transactions on Information Theory*, *38*(2), 617–643.

Mallat, S., & Peyré, G. (2007). A review of bandlet methods for geometrical image representation. *Numerical Algorithms*, *44*(3), 205–234.

Mallat, S., & Zhong, S. (1992). Characterization of signals from multiscale edges. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *14*(7), 710–732.

Matas, J., Chum, O., Urban, M., & Pajdla, T. (2002). Robust wide baseline stereo from maximally stable extremal regions. In *British machine vision conference* (pp. 384–393).

Mikolajczyk, K., & Schmid, C. (2004). Scale and affine invariant interest point detectors. *International Journal of Computer Vision*, *60*(1), 63–86.

Mikolajczyk, K., & Schmid, C. (2005). A performance evaluation of local descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *27*(10), 1615–1630.

Mikolajczyk, K., Tuytelaars, T., Schmid, C., Zisserman, A., Matas, J., Schaffalitzky, F., Kadir, T., & Gool, L. V. (2005). A comparison of affine region detectors. *International Journal of Computer Vision*, *62*(1), 43–72.

Rudin, W. (1991). *Functional analysis*. New York: McGraw-Hill.

Tuytelaars, T., & Gool, L. V. (2004). Matching widely separated views based on affine invariant regions. *International Journal of Computer Vision*, *59*(1), 61–85.

Witkin, A. (1983). Scale-space filtering. In *Proceedings of the eight international joint conference on artificial intelligence*.