

Highlight on a feature extracted at fine scales: the pointwise Lipschitz regularity

C. Damerval¹ and S. Meignen²

¹ Dept. of Computer Science, Katholieke Universiteit Leuven (Belgium)

² Laboratoire Jean Kuntzmann (LJK), University of Grenoble (France)

Abstract. The aim of this paper is to study the robustness of the pointwise Lipschitz regularity in 2D, which is a measure of the local regularity of the intensity function associated to an image. This regularity can be efficiently computed by an approach based on fine scales. We assess its robustness when the image undergoes various transformations, especially geometric ones. The results we obtain show that the pointwise Lipschitz regularity is a suitable feature for applications in computer vision.

Keywords: Lipschitz regularity, invariance properties, wavelet decompositions, multiscale edge detection, extraction of characteristic values, robustness to 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 transformations such as scale change, multiscale approaches were put forward. In particular, methods based on the *Scale-Space* theory [1–3] turned out as successful in computer vision. These can evidence regions of interest, which are stable through local geometric deformations [4]. These regions are identified by their location and characteristic scale, and their content can be quantified by a suitable descriptor [5]. Recent works compared state-of-the-art interest regions detectors [6] and region descriptors [7]. Existing methods proved to be efficient for one type of scene or transformation [8–10]; however, no method outperforms the others in all cases, so combining different kinds of features seems relevant.

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). It was widely studied in 1D, especially in the case of multifractal signals [11, 12], and also applied to the characterization of singularities [13, 14] and landmark registration [15]. In 2D, methods based on regularity measures were also put forward with applications to textured images [16, 17], the regularity being used from a global point of view. Besides, recent advances in edge detection using multiscale approaches [18] were used for object detection [19]. New developments using the multiscale SIFT descriptor were also recently proposed [20].

Here we are concerned with the pointwise regularity in 2D. More precisely, a multiscale approach focused on fine scales allows to compute numerically values of regularity α . As we will see, this regularity appears as a relevant feature to detect points of interest and could be profitably used as an image descriptor: on the one hand it has invariance properties (especially regarding geometric transformations), and on the other hand values of the pointwise Lipschitz regularity α can be efficiently computed.

This paper is organized as follows. We first present the notion of regularity α in 2D, its invariance properties. So as to numerically estimate α in 2D, we recall an algorithm based on a multiscale edge detector [21], which gives pointwise estimations of α at edge points of the image. We also present a methodology so as to compare values of α between two images related by a geometric deformation. Finally, an evaluation procedure allows to assess the robustness of the regularity α , for natural scenes viewed under various imaging conditions. The obtained results show that the regularity α makes up a robust feature.

2 Regularity α in the context of image analysis

We consider an image (in level of gray) given by its intensity function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$. We first present briefly the definition of Lipschitz regularity in 2D, inferred from the 1D definition [13]. This leads to the notion of regularity α . Then we recall a known algorithm for computing the values of the regularity α . We also investigate its invariance properties, especially from a practical point of view.

2.1 Notion of regularity α in 2D – Invariance properties

The Lipschitz regularity generalizes the usual notion of regularity.

Definition 1. (*1D Lipschitz regularity*) Given $\alpha \in]0, 1[$, a function $f : \mathbb{R} \rightarrow \mathbb{R}$ is α -Lipschitz at $x_0 \in \mathbb{R}$ if there exists a neighborhood V of x_0 and $A > 0$ so that

$$\forall x \in V, |f(x) - f(x_0)| \leq A|x - x_0|^\alpha$$

This can be extended for $\alpha \in \mathbb{R}$. In particular, for $\alpha = n \in \mathbb{N}^*$, this corresponds to a locally C^n function. Besides, for $\alpha < 0$, this definition can be generalized thanks to the theory of distributions (see details in [14]).

Definition 2. (*2D Lipschitz regularity*) Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ and $x_0 \in \mathbb{R}^2$. For $\theta \in [0, 2\pi[$, we define $f_\theta : \mathbb{R}_+^* \rightarrow \mathbb{R}$ 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

$$\forall \theta \in [0, 2\pi[, f_\theta \text{ } \alpha\text{-Lipschitz at } 0$$

Note that this definition agrees with the usual definition of the Lipschitz regularity. Indeed, when $\alpha \in]0, 1[$, provided f is α -Lipschitz at x_0 , we can write

$$\begin{aligned} |f(x) - f(x_0)| &= |f(x_0 + (h \cos \theta, h \sin \theta)) - f(x_0)| = |f_\theta(h) - f_\theta(0)| \leq Ah^\alpha \\ |f(x) - f(x_0)| &\leq A||x - x_0||^\alpha \text{ with } A > 0, \text{ for } x \text{ in a neighborhood of } x_0. \end{aligned}$$

Definition 3. (*Regularity α*) Let $f : \mathbb{R}^2 \longrightarrow \mathbb{R}$ and $x_0 \in \mathbb{R}^2$. The regularity α of f at x_0 is defined as

$$\alpha = \alpha(f, x_0) = \inf\{\alpha_0 \in \mathbb{R}, f \text{ } \alpha_0\text{-Lipschitz at } x_0\} \quad (1)$$

The relevance of α arises from its invariance properties: the regularity α appears as a characteristic value. In particular, let us study the case of a constant affine deformation (so including rotation and scale change), widely studied in the Scale-Space theory [2].

Proposition 1. (*Influence of an affine deformation on the regularity α*)
Let $f : \mathbb{R}^2 \longrightarrow \mathbb{R}$, and g defined by

$$\forall x \in \mathbb{R}^2, g(x) = f(Bx), \text{ with } B \text{ a } 2 \times 2 \text{ invertible matrix} \quad (2)$$

Then, for $\alpha \in]0, 1[$, $\alpha(f, x_0) = \alpha(g, y_0)$ with $y_0 = B^{-1}x_0$.

Proof. According to def. 3, there exists $A > 0$ so that

$$\forall \theta \in [0, \pi[, |f(x_0) - f(x_0 + hu_\theta)| \leq Ah^\alpha \quad (3)$$

Let us study the regularity of g at $y_0 = B^{-1}x_0$. For $\theta \in [0, 2\pi[$ we have

$$|g(y_0) - g(y_0 + hu_\theta)| = |g(B^{-1}x_0) - g(B^{-1}x_0 + hu_\theta)| \quad (4)$$

$$= |f(x_0) - f(x_0 + hBu_\theta)| \quad (5)$$

Moreover, since $Bu_\theta = \lambda u_{\theta'}$ with $\lambda \in \mathbb{R}^*$ and $\theta' \in [0, 2\pi[$:

$$|g(y_0) - g(y_0 + hu_\theta)| = |f(x_0) - f(x_0 + h\lambda u_{\theta'})| \quad (6)$$

$$\leq (A|\lambda|^\alpha)h^\alpha \quad (7)$$

So there exists $A' > 0$ so that

$$\forall \theta \in [0, 2\pi[, |g(y_0) - g(y_0 + hu_\theta)| \leq A'h^\alpha \quad (8)$$

and g is α -Lipschitz at x_0 .

Then, 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. \square

Note this invariance property may not always hold in practice. Indeed, when B becomes nearly singular (case of extreme deformations), λ can be very small when u_θ is an eigenvector of B . So there may be numerical instabilities for extreme deformations. However, as we will see, the regularity α yields a significant robustness for wide-ranging transformations (and not only small deformations).

Now, let us discuss more general transformations, given by

$$\forall x \in \mathbb{R}^2, g(x) = f(v(x)) \text{ with } v : \mathbb{R}^2 \longrightarrow \mathbb{R}^2 \quad (9)$$

In this general context, note that the regularity α is not necessarily preserved: depending on the regularity of v , g may be more regular than f , resulting in a higher regularity α for g than for f . Nevertheless, we point out that it can be preserved in many practical cases, especially when considering image edges [22]. Note that in the case of an image representing an edge, f is regular along the tangent to the edge and irregular along the normal direction – see Fig.1(a). In this regard, 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 further explain how to compute this direction and estimate α (section 2.2). More generally, since transformations such as local affine deformations do not alter the topology of the edges, the regularity α on these edges should be preserved – see Fig.1(b).

2.2 Numerical computation through a multiscale approach

So as to compute numerical values of α , we use a known approach based on a multiscale edge detector [21]. Let us recall briefly some aspects of this detector. According to Canny [22], edge points correspond to locations where the magnitude of the gradient attains a local maximum in the direction of the gradient – which is the direction of maximum irregularity. A generalization of Canny’s detector was put forward by Mallat, using wavelet decompositions [21]. This allows to detect edge points, and also to compute an accurate estimation of the regularity α at these edge points. This computation of α is carried out using a linear regression at the finest scales. Besides, denoting N the size of the data ($N = n^2$ for an image $n \times n$), this formulation can be computed in $O(N)$, thus allowing a fast computation. In summary, this detector is known as an efficient method so as to compute numerical values of α . We emphasize that this method is focused on the finest scales, and that it gives pointwise estimations of the regularity α . Given an image f , the output of this detector can be expressed as a set

$$\{(x_i, y_i, \alpha_i) \in \mathbb{R}^3, 1 \leq i \leq n_f\} \quad (10)$$

where n_f is the number of detected edge points (x_i, y_i) , each being associated to a value α_i . Since edge points correspond to singularities (where f may not be differentiable), the obtained values can be negative: typically a boundary leads to $\alpha = 0$, a line to $\alpha = -1$, and an isolated point to $\alpha = -2$. For natural images, we obtain various values; for instance, given the image represented on Fig.2(a), we represent the density associated to the regularity α of detected edge points on Fig.2(c). In this case, 95% of the computed values are within $[-1.4, 0.8]$. Besides, some parameters of the detector (like thresholding) allow to tune the number of edge points. We use here a light thresholding in our numerical experiments, so as to obtain a large number of values of regularity α . This is consistent insofar as we want to evaluate the robustness of the regularity α from a practical point of view.

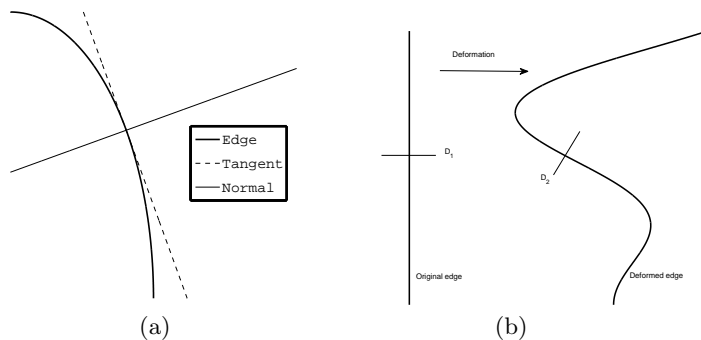


Fig. 1. (a) At a point belonging to an edge line, the Lipschitz regularity is minimal along the normal direction; along this direction, the regularity α can be accurately computed. (b) If an edge undergoes a deformation which does not change its topology, the regularity α is preserved (D_1, D_2 : directions of maximum irregularity.)

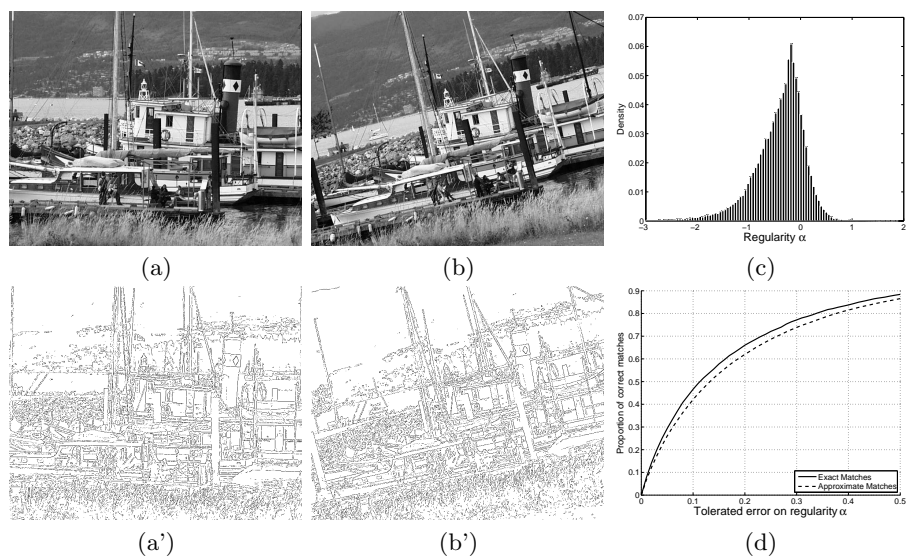


Fig. 2. (a,b) Two images related by a known affine deformation; (a', b') Detected edge points; (c) Density of the regularity α , based on edge points represented on (a'); (d) Errors on the estimation of α for exact matches (EM) and approximate matches (AM) between edge points of (a') and (b').

2.3 Empirical 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 edge points. For that purpose we consider an original image X_0 (see Fig.2(a)) and a deformed image X_1 (see Fig.2(b)) related by a known affine deformation. This homography allows to carry out point-to-point correspondences and thus to compare the values of regularity between the two images (see Fig.2(a',b')). Given X_0 and X_1 , the detector leads to two sets of edge points with associated values of regularity α :

$$S_0 = (x_i^0, y_i^0, \alpha_i^0)_{i \in I_0} \text{ and } S_1 = (x_j^1, y_j^1, \alpha_j^1)_{j \in I_1}$$

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 between the sets S_1 and S_0 . At this step, we have two possible choices: either exact matches (EM) for which a projected point of S_1 corresponds exactly to a point of S_0 ; or approximate matches (AM), by tolerating an error of 1 pixel. 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 (for one matched pair) as

$$d_\alpha = d_\alpha(i, j) = |\alpha_i^0 - \alpha_j^1| \quad (11)$$

Finally, a match is said correct if $d_\alpha < \epsilon$, where the parameter ϵ is a tolerated error on α . We point out there is a certain freedom of choice for ϵ , which should depend on the application. In this regard, there is a trade-off between too low values (refusing any numerical error on α) and too high values of ϵ (not taking into account α). Let us now study the effect of this parameter $\epsilon > 0$, by comparing the sets S_0 and S_1 : we represent on Fig.2(d) the proportion of correct matches depending on ϵ , for both EM and AM. As expected, it increases with respect to the parameter ϵ . More precisely, a good result lies in the fact that it increases rapidly for small values ϵ , becoming thus significant: indeed, when ϵ exceeds 0.3, this proportion attains almost 80%. Moreover we observe that the results are only slightly better for EM than AM, so that it can be relevant to consider AM to define a descriptor since the number of extracted points is significantly larger. So these first results show that the regularity α estimated at edge points is a feature robust to affine deformations. Let us now evaluate the robustness of this feature in a more general context, when various transformations are applied to natural images.

3 Quantifying the robustness of the regularity α

3.1 Evaluation procedure

We consider 8 sequences, each consisting of 6 images $(X_k)_{0 \leq k \leq 5}$: ZoomRotation1, ZoomRotation2, Viewpoint1, Viewpoint2, Blur1, Blur2, Jpeg and Light (see Fig.3). For each sequence, the 6 images represent a given scene viewed under a certain imaging condition. For instance, considering the sequence Viewpoint1

(see Fig.4), each image X_k ($1 \leq k \leq 5$) corresponds to a change of viewpoint applied to the reference image X_0 . The relevance of these sequences lies in different aspects. First they represent various objects: textured scenes – repeated textures, see Fig.3(a,d) – and structured ones – homogeneous regions with edges boundaries, see Fig.3(b,c). Secondly the imaging conditions are wide-ranging: geometric deformations and specific transformations like JPEG compression. Thirdly the degree of these transformation can be significant (scale change up to 4, angle of viewpoint up to 60° , JPEG compression rate up to 98%). Finally, we mention the sequences ZoomRotation1, ZoomRotation2, Viewpoint1 and Viewpoint2 correspond to actual camera displacements; the sequences Blur1, Blur2 and Light correspond to camera operations (varying the camera focus or shutter speed); for the sequence Jpeg, different levels of JPEG compression were obtained by a software. For illustration purposes, we represent on Fig.3 the images X_0 and X_5 associated to every sequence. For more details, see <http://www.robots.ox.ac.uk/~vgg/research/affine>.

For a given set of images $(X_k)_{0 \leq k \leq 5}$ associated to a sequence (viewpoint change for instance, see Fig.4), we carry out the following procedure:

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

$$\mathcal{C}^k = \{ (p_i^0, p_j^k) \text{ matched according to a geometric criterion} \} \quad (12)$$

This leads to a certain number of correspondences (NC) $\#\mathcal{C}^k$.

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

$$\mathcal{C}_\epsilon^k = \{ (p_i^0, p_j^k) \in \mathcal{C}^k, d_\alpha = |\alpha_i^0 - \alpha_j^k| < \epsilon \} \quad (13)$$

and compute the matching score, representing the proportion of correct matches:

$$S_k = \frac{\#\mathcal{C}_\epsilon^k}{\#\mathcal{C}^k} \quad (14)$$

This score reflects the robustness of the regularity α . We will study its evolution with respect to ϵ for the sequence Viewpoint1 (Fig.5), and also with respect to k for all sequences (Fig.6, for fixed ϵ). Note that in step 2, the matches based on a geometric criterion can be either exact or approximate (as described in section 2.3); we study both EM and AM. In step 3, one may have to choose the parameter ϵ , representing a tolerated error on α (as seen in section 2.3). In our experiments we use $\epsilon = 0.3$, for which a high proportion of the matches (almost 80%) are deemed correct in the case of the affine deformation studied in section 2.3 (see Fig.2(d)). Note also that this choice allows to identify clearly boundaries ($\alpha = 0$), lines ($\alpha = -1$) and isolated points ($\alpha = -2$).




Fig. 3. Sample of data set, representing X_0 (reference image) and X_5 (highest degree of transformation) associated to: (a, b) Scale change and rotation; (c, d) Viewpoint change; (e, f) Blur; (g) JPEG compression; (h) Illumination change.




Fig. 4. (Top) Complete sequence Viewpoint1 (viewpoint change): 6 images X_0, \dots, X_5 ; (Bottom) Associated edge points.

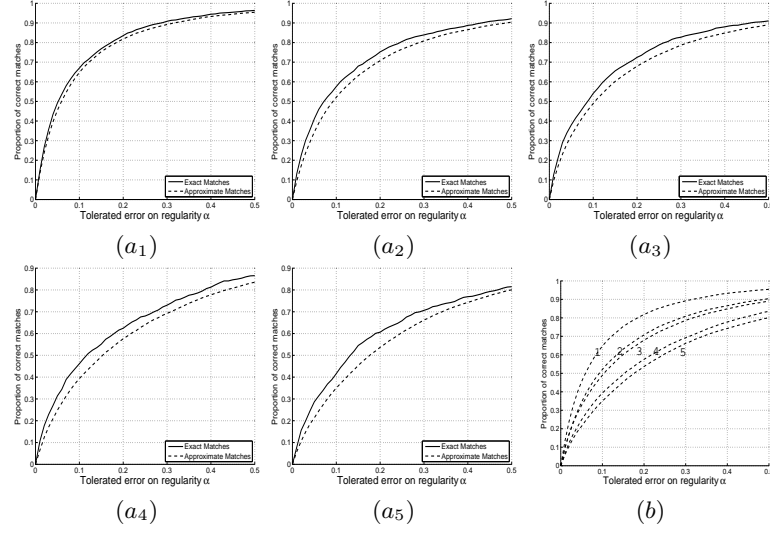


Fig. 5. Proportion of correct matches with respect to the tolerated error on α , between the images X_0 and X_k ($1 \leq k \leq 5$) associated to the sequence Viewpoint1 (see Fig.4). (a_k) Comparison between X_0 and X_k , for both EM and AM; (b) Comparison between X_0 and all X_k , for AM (the curves 1–5 correspond to those of fig.(a1–a5) for AM).

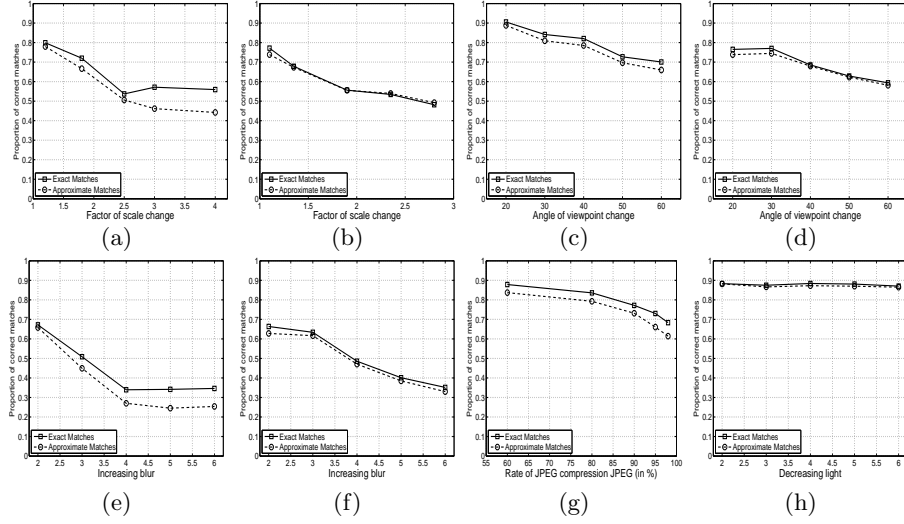


Fig. 6. (a–h) Matching scores ($\epsilon = 0.3$) associated to the sequences of Fig.3(a–h).

3.2 Analysis of the results

Study of the sequence Viewpoint1 (Fig.5)

First, we consider the sequence Viewpoint1, associated to increasing angles of viewpoint changes ($k = 1 : 20^\circ, k = 2 : 30^\circ, k = 3 : 40^\circ, k = 4 : 50^\circ, k = 5 : 60^\circ$). For each k ($1 \leq k \leq 5$), we represent on Fig.5(a_k) the evolution of S_k with respect to the tolerated error on α (parameter ϵ), for both EM and AM. The obtained results are similar to the affine deformation studied in section 2.3; indeed, EM yield slightly better than AM in all cases, and the proportion of correct matches increases with the tolerated error on α (parameter ϵ). In particular, for $\epsilon = 0.3$, we note that the proportion of correct matches (for EM and AM) exceeds: 80% for the angles $20^\circ, 30^\circ$ and 40° , see Fig.5(a_1 – a_3); 60% for the angles 50° and 60° , see Fig.5(a_4 – a_5). Besides, we display on Fig.5(b) all the preceding curves (only for EM), representing the degradation of the estimation of regularity α as the degree of the viewpoint change increases. Note that as soon as ϵ is larger than 0.2, the proportion of correct matches exceeds 50%, even for significant changes of viewpoint. This good result shows the regularity α makes up a robust feature.

Study of all sequences (Fig.6)

Now, considering the 8 sequences, for both approximate and exact matches, we represent on Fig.6 the curves (k, S_k). More precisely, each graph of Fig.6 describes the performance for one particular sequence associated to a certain image transformation. For instance, Fig.6(c) refers to the sequence Viewpoint1 (associated to a viewpoint change), represented partially on Fig.3(c) (and comprehensively on Fig.4). This allows to assess the robustness of the estimation of the regularity α in general situations, which is the main objective of this paper. We emphasize that a method is all the better than it leads to higher scores and that these scores are stable, i.e, they remain high when the degree of deformation increases. On the basis of the matching scores (see section 3.1), we can evaluate how robust the estimation of α is, under various imaging conditions (Fig.6). 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 degree 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 section 2.3, see Fig.2). We do not observe a significant difference between textured (Fig.6(a,d)) and structured scenes (Fig.6(b,c)); however, the structured scene of Fig.6(c) performs better, due to the presence of clear edges. Let us now detail the analysis of the results for each transformation.

Scale change and rotation, Fig.6(a,b): for both textured and structured scenes, 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, Fig.6(c,d): the performance decreases moderately, remaining high for the structured scene (between 0.9 and 0.7, Fig.6(c)) and good for the textured scene (between 0.8 and 0.6, Fig.6(d)). In this regard, note that both sequences Viewpoint1 and Viewpoint2 contain distinct edge, which are moderately affected by a viewpoint change.

Blur, Fig.6(e,f): the performance decreases rapidly (from 0.7 to 0.3) for both structured scenes and textured ones. It is not surprising to obtain these average results since the blur modifies the regularity α (edges are smoothed). It is well known that a smoothing operation alters the regularity α ; yet, one can retrieve the regularity α when the smoothing kernel is known [14].

JPEG compression, Fig.6(g): the performance remains high (between 0.9 and 0.6), decreasing steadily. So JPEG artifacts have little impact on the regularity α , even if JPEG compression tends to blur sharp edges.

Light change, Fig.6(h): very high performance (stable, close to 0.9). Since illumination change does not alter the structure of the edges, the regularity α is not affected by such changes.

In conclusion, 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 degree of transformation. This emphasizes the fact that the regularity α is characteristic of the kind of edge. In a blurry context, the computation of α seems less reliable; nevertheless this can be improved by using higher scales in the detector (provided the edges are sufficiently far apart). 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: the scores are only slightly inferior for AM than for EM, while AM leads to a greater number of points (which is important for image description). Note also that in an application such as image matching, increasing the parameter ϵ leads to a limited number of matched pairs (for which the computed values are closer, see section 2.3).

3.3 Discussion

Let us discuss now some aspects of the regularity α . First, concerning the way of computing values of α , note there exist methods based on wavelet transforms that allow to estimate the regularity α at any point of the image [23]. Here we focused on edges, since they appear to be robust to various transformations of the image. Moreover, so as to improve the matching performance, the number of detected edge points can be limited. This can be done by selecting only the highest responses (threshold on the modulus of the gradient), or by considering higher scales. This will result in evidencing only the most salient edges. Here, we considered all the values for regularity α , showing that a significant proportion of the computed values α corresponds to values for which the estimation is robust.

In addition, we can compare certain aspects of our approach with works related to interest regions [6, 7], to the extent they allow to characterize some objects present in the image. To that regard, it is important to note that we focus on position and pointwise regularity α , whereas these methods are based on position and characteristic scale. On the one hand, such methods use this characteristic scale so as to define interest regions, and then compute associated descriptors

which characterize their content. On the other hand, our method gives pointwise features, so there is no region of interest associated to the Lipschitz regularity (the definition of such regions seems possible, but is not straightforward).

Besides, we can compare the performance measures of these two methods. Since the criteria used are different, we can only draw conclusions from the shape of the curves. Numerically we observe that our method is more stable than those based on interest regions: 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 for ZoomRotation1 and Blur1; equivalent for Light, Blur2. So, compared to the best state-of-the-art methods, the regularity α (computed at edge points) yields a significant robustness to various image transformations.

4 Conclusions and perspectives

In this paper, we studied the regularity α in the context of interest point detection, focusing on edge points. This approach is based on fine scales (pointwise features) which differs from other methods based on coarser scales (local features). We explained why certain transformations of the image do not change the pointwise regularity α at such edge points. Hence the regularity α stands out as a characteristic value.

The main contribution of our work lies in quantifying the robustness of the 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. This leads to good results of robustness concerning geometric deformations – such as viewpoint change, scale change and rotation – and also JPEG compression and illumination change. So the regularity α (computed at edges) appears as a relevant feature for various tasks in computer vision.

In terms of perspectives, let us point out potential applications of the regularity α . A first application may consist on clustering edge points into edges (1D-curves) since these have connexity properties. Instead of relying only on a distance measure, such a clustering would use a criterion based on both distance and regularity α . Secondly, the regularity α could be used complementary to interest region descriptors: the estimated regularity α at all edge points within a given region may help to characterize the content of the region. For that purpose, it is interesting that our method can evidence a great number of features. Besides, the regularity α has potential applications to image registration, in particular feature-based methods (thanks to the identification of lines, curves, points and corners). Eventually, the regularity α appears as an interesting additional feature, additional to other existing local features: an integration of the pointwise Lipschitz regularity in existing detectors will certainly improve their performance.

Acknowledgement We would like to thank the referees for relevant suggestions, and also Prof. M. Jansen from the K.U. Leuven for useful comments.

References

1. Iijima, T.: Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bull. of the Electrotechnical Laboratory* **26** (1962) 368–388
2. Lindeberg, T.: *Scale Space Theory in Computer Vision*. Kluwer (1994)
3. Witkin, A.: Scale-space filtering. *Proceedings of the 8th International Joint Conference on Artificial Intelligence* (1983) 1019–1021
4. Lindeberg, T.: Feature detection with automatic scale selection. *International Journal of Computer Vision* **30**(2) (1998) 77–116
5. Lowe, D.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60**(2) (2004) 91–110
6. Mikolajczyk, K., Tuytelaars, T., Schmid, C., Zisserman, A., Matas, J., Schaffalitzky, F., Kadir, T., Gool, L.V.: A comparison of affine region detectors. *International Journal of Computer Vision* **62**(1) (2005) 43–72
7. Mikolajczyk, K., Schmid, C.: A performance evaluation of local descriptors. *IEEE Transactions on PAMI* **27**(10) (2005) 1615–1630
8. Kadir, T., Zisserman, A., Brady, M.: An affine invariant salient region detector. *Proceedings of the 8th European Conference on Computer Vision* (2004) 228–241
9. Matas, J., Chum, O., Urban, M., Pajdla, T.: Robust wide baseline stereo from maximally stable extremal regions. *British Machine Vision Conf.* (2002) 384–393
10. Tuytelaars, T., Gool, L.V.: Matching widely separated views based on affine invariant regions. *International Journal of Computer Vision* **59**(1) (2004) 61–85
11. Arneodo, A., Bacry, E., Jaffard, S., Muzy, J.F.: Singularity spectrum of multifractal functions involving oscillating singularities. *Journal of Fourier Analysis and Applications* **4**(2) (1998) 159–174
12. Benassi, A., Cohen, S., Istas, J., Jaffard, S.: Identification of filtered white noises. *Stochastic Processes and their Applications* **75**(1) (1998) 31–49
13. Jaffard, S., Meyer, Y.: Wavelet methods for pointwise regularity and local oscillations of functions. *American Mathematical Society* (1996)
14. Mallat, S., Hwang, W.L.: Singularity detection and processing with wavelets. *IEEE Transactions on Information Theory* **38**(2) (1992) 617–643
15. Bigot, J.: Automatic landmark registration of 1d curves. In: *Recent advances and trends in nonparametric statistics*. Elsevier (2003) 479–496
16. Deguy, S., Debain, C., Benassi, A.: Classification of texture images using multiscale statistical estimators of fractal parameters. *British Machine Vision Conference* (2000) 192–201
17. Kaplan, L.M., Kuo, C.C.: Texture roughness analysis and synthesis via extended self-similar model. *IEEE Transactions on PAMI* **17**(11) (1995) 1043–1056
18. Martin, D., Fowlkes, C., Malik, J.: Learning to detect natural image boundaries using local brightness, color and texture cues. *IEEE Transactions on PAMI* **26**(5) (2004) 530–549
19. Ferrari, V., Fevrier, L., Jurie, F., Schmid, C.: Groups of adjacent contour segments for object detection. *IEEE Transactions on PAMI* **30**(1) (2008) 36–51
20. Brown, M., Lowe, D.: Automatic panoramic image stitching using invariant features. *International Journal of Computer Vision* **74**(1) (2007) 59–73
21. Mallat, S., Zhong, S.: Characterization of signals from multiscale edges. *IEEE Transactions on PAMI* **14**(7) (1992) 710–732
22. Canny, J.: A computational approach to edge detection. *IEEE Transactions on PAMI* **8**(6) (1986) 679–698
23. Mallat, S.: *A wavelet tour of signal processing*. Academic Press (1998)