

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2017.DOI

Performance Characterization of Image Feature Detectors in Relation to the Scene Content Utilizing a Large Image Database

BRUNO FERRARINI¹, SHOAB EHSAN¹, ALEŠ LEONARDIS², NAVEED UR REHMAN³, and KLAUS D. MCDONALD-MAIER¹

¹University of Essex, School of Computer Science and Electronic Engineering, Colchester CO4 3SQ, UK

²School of Computer Science, University of Birmingham, Birmingham, UK

³Dep. of Electrical Engineering, COMSATS Institute of Information Technology, Islamabad, Pakistan

Corresponding author: Bruno Ferrarini (e-mail: bferra@essex.ac.uk).

This work has been supported by the UK Engineering and Physical Sciences Research Council EPSRC [EP/K004638/1, EP/R02572X/1 and EP/P017487/1]

ABSTRACT Selecting the most suitable local invariant feature detector for a particular application has rendered the task of evaluating feature detectors a critical issue in vision research. Although the literature offers a variety of comparison works focusing on performance evaluation of image feature detectors under several types of image transformations, the influence of the scene content on the performance of local feature detectors has received little attention so far. This paper aims to bridge this gap with a new framework for determining the type of scenes which maximize and minimize the performance of detectors in terms of repeatability rate. The results are presented for several state-of-the-art feature detectors that have been obtained using a large image database of 20482 images under JPEG compression, uniform light and blur changes with 539 different scenes captured from real-world scenarios. These results provide new insights into the behavior of feature detectors.

INDEX TERMS Feature extraction, Image analysis, Performance analysis, Feature Detector, Comparison, Repeatability.

I. INTRODUCTION

Local feature detection has been a challenging area of interest for the computer vision community for some time. A large number of different approaches have been proposed so far, thus making evaluation of image feature detectors an active research topic in the last decade or so. Most evaluations available in the literature focus mainly on characterizing feature detectors' performance under different image transformations without analyzing the effects of the scene content in detail. In [1], the feature tracking capabilities of some corner detectors are assessed utilizing static image sequences of a few different scenes. Although the results permit to infer a dependency of the detectors' performance on the scene content, the methodology followed is not specifically intended to highlight and formalize such a relationship, as no classification is assigned to the scenes. The comparison work in [2] gives a formal definition for textured and structured scenes and shows the repeatability rates of six feature detec-

tors. The results provided by [2] show that the content of the scenes influences the repeatability but the framework utilized and the small number of scenes included in the datasets [3] do not provide a comprehensive insight into the behavior of the feature detectors with different types of scenes. In [4], the scenes are classified by the complexity of their 3D structures in complex and planar categories. The repeatability results reveal how detectors perform for those two categories. The limit in the generality of the analysis done in [4] is due to the small number and variety of the scenes employed, whose content are mostly human-made. This paper aims to help better understand the effect of the scene content on the performance of several state-of-the-art local feature detectors. The main goal of this work is to identify the biases of these detectors towards particular types of scenes, and how those biases are affected by three different types and amounts of transformations (JPEG compression, blur and uniform light changes). The methodology proposed utilizes

the improved repeatability criterion presented in [5], as a measure of the performance of feature detectors, and the large database [6] of images consisting of 539 different real-world scenes containing a wide variety of different elements. This paper offers a more complete understanding of the evaluation framework first described in the conference version [7] and further developed in [8], providing additional background, description, insight, analysis and evaluation.

The remainder of the paper is organized as follows. Section II provides an overview of the related work in the field of feature detector evaluation and scene taxonomy. In Section III, the proposed evaluation framework is described in detail. Section IV is dedicated to the description of the image database utilized for the experiments. The results utilizing the proposed framework are presented and discussed in Section V. Finally, Section VI presents the conclusions.

II. RELATED WORK

The contributions to the evaluation of local feature detectors are numerous and vary based on: 1) the metric used for quantifying the detector performance, 2) the framework/methodology followed and 3) the image databases employed. Repeatability is a desirable property for feature detectors as it measures the grade of independence of the feature detector from changes in the imaging conditions. For this reason, it is frequently used as a measure of performance of local feature detectors. A definition of repeatability is given in [9] where, together with the information content, it is utilized as a metric for comparing six feature detectors. A refinement of the definition of repeatability is given in [10], and used for assessing six state-of-the-art feature detectors in [2] under several types of transformations on textured and structured scenes. Two criteria for improved repeatability measure are introduced in [5] that provide results which are more consistent with the actual performance of several popular feature detectors on the widely-used Oxford datasets [3]. The improved repeatability criteria are employed in the evaluation framework proposed in [11] and in [12], which presents a method to assess the performance bounds of detectors.

Moreover, repeatability is used as a metric for performance evaluation in [13] and [4] that utilize non-planar, complex and simple scenes.

The performance of feature detectors has also been assessed employing metrics other than repeatability. The performance measure in [14] is completeness, while feature coverage is used as a metric in [15]. The feature detectors have also been evaluated in the context of specific applications, such as in [1], where corner detectors are assessed in the context of point feature tracking applications.

III. THE PROPOSED EVALUATION FRAMEWORK

The proposed framework has been designed by keeping in mind the objective of evaluating the influence of scene content on the performance of a wide variety of state-of-the-art feature detectors. A proper application of such a framework

requires a large image database I organized in a series of n datasets. Each dataset needs to contain images from a single scene with different amounts of image transformation. The images included in such a database should be taken from a large variety of different real-world scenarios. The proposed framework consists of the steps discussed below.

A. REPEATABILITY DATA

The framework is based on the repeatability criterion described in [5], whose consistency with the actual performance of a wide variety of feature detectors has been proven across well-established datasets [3]. As proposed in [5], the repeatability rate is defined as follows:

$$\text{Repeatability} = \frac{N_{rep}}{N_{ref}}, \quad (1)$$

where N_{rep} is the total number of repeated features and N_{ref} is the number of interest points in the common part of the reference image. Let A and P be the sets of indices representing the m discrete amount of transformation and the scenes respectively.

$$A = \{1, 2, 3, \dots, m\}, \quad (2)$$

$$P = \{1, 2, 3, \dots, n\}, \quad (3)$$

where m corresponds to the maximum amount of transformation and 1 relates to the reference image (zero transformation); n is the total number of scenes and each scene is utilized to build one separate dataset, thus finally resulting in n datasets in total. Let B_{kd} be the set of repeatability rates computed for step k (corresponding to k image transformation amount) for a feature detector d across n datasets (which implies repeatability values for n scenes):

$$B_{kd} = \{B_{1kd}, B_{2kd}, \dots, B_{nkd}\}. \quad (4)$$

Each set B_{kd} contains n repeatability ratios, one for each dataset. In particular, for the image database utilized in this work for the experiments [6], n is 539 while maximum value of k is 10 or 14 depending on which transformation is considered. Thus, B_{kd} includes 539 values of repeatability for the k^{th} step.

B. SCENE RANKINGS

The top and lowest rankings for each detector d are built selecting the j highest and the lowest repeatability scores at k amount of image transformation. Let $T_{kd}(j)$ and $W_{kd}(j)$ the sets containing the indices of the scenes whose repeatability falls in the top and lowest ranking respectively:

$$T_{kd}(j) = \{S_{kd(1)}, S_{kd(1)}, \dots, S_{kd(j)}\}, \quad (5)$$

$$W_{kd}(j) = \{S_{kd(n)}, S_{kd(n-1)}, \dots, S_{kd(n-j+1)}\}, \quad (6)$$

where $S_{kd(i)} \in P$ is the scene index corresponding to the i^{th} highest repeatability score obtained by the detector d for the scene under k amount of transformation. Thus,




FIGURE 1: The reference image of three scenes and the effect of the application of 60% of light reduction, 98% of JPEG compression rate and 4.5σ Gaussian blur.

in accordance with this notation, $S_{kd(1)}$ is the scene for which the detector scored the best repeatability score, $S_{kd(2)}$ corresponds to the second highest repeatability rate, $S_{kd(3)}$ to the third highest and so on, until $S_{kd(n)}$ which is for the lowest one.

TABLE 1: Classification labels and criteria

Location Type	Outdoor	Indoor scene and close-up a single or of a few objects.
	Indoor	The complement of above.
Object Type	Human-made	Elements are mostly artificial.
	Natural	Elements are mostly natural.
Complexity	Simple	A few edges with quite regular shapes.
	Complex	A large number of edges with fractal-like shapes.

C. SCENE CLASSIFICATION

The scenes are attributed with three labels on the basis of human judgment. As described in Table 1, each label is dedicated to a particular property of the scene and has been assigned independently from the others. These attributes are: the location type (f), which may take the label outdoor or indoor, the type of the elements contained (g), which may take the label natural or human-made, and the perceived complexity of the scene (h), which may take the label simple or complex. Fig. 2 shows a sample of the scenes from the image database [6] utilized for the experiments grouped so that each row shows scenes sharing the same value for one of the three labels f , g and h . Scene 9 is tagged as outdoor and, along with scene 76 and 17, contains natural elements. The scenes 40, 530 and 373 are labeled as human-made and the first is also classified as indoor. The scene 530 is categorized as a simple scene as it includes a few edges delimiting well contrasted areas. Although the main structures (broccolis borders) can be identified in scene 76, the rough surface of the broccolis is information rich that results in labeling this scene as complex.

D. RANKING TRAIT INDICES

The labels of the scenes included in the rankings, (5) and (6), are examined in order to determine the dominant types of scenes. For each ranking $T_{kd}(j)$ and $W_{kd}(j)$, the ratios of scenes classified as *outdoor*, *human-made* and *simple* are computed. Thus, three ratios are associated to each ranking where higher values mean higher share of the scene type associated:

$$\forall S_i \in T_{kd} : T_{kd} \cdot [F, G, H] = \frac{\sum S_i \cdot [f, g, h]}{j}, \quad (7)$$

$$\forall S_i \in W_{kd} : W_{kd} \cdot [F, G, H] = \frac{\sum S_i \cdot [f, g, h]}{j}. \quad (8)$$

These vectors contain three measures which represent the extent of the bias of detectors. For example, if a top ranking presents $F = 0.1$, $G = 0.25$ and $H = 0.8$, it can be concluded that the detector, for the given amount of image



FIGURE 2: Some images from the database utilized for the experiments. Each row shows images belonging to the same category: outdoor, non-outdoor, human-made content, natural content, simple and complex edges.

transformation, works better with scenes where its' element are mostly natural (low G), with simple edges (high H) and that are not outdoor (low F). As opposed to that, if the same indices were for a lowest ranking it could be concluded that the detector obtains its lowest results for non-outdoor (F) and natural (G) scenes with low edge complexity (H).

IV. IMAGE DATASET

The image database used for the experiments is discussed in this section and is available at [6]. It contains a large number of images, 20482, from real-world scenes. This database includes a wide variety of outdoor and indoor scenes showing both natural and human-made elements. The images are organized in three groups of datasets, corresponding to the three transformations: JPEG compression, uniform light and Gaussian blur changes. Each dataset includes a reference image for the particular scene and several images of the same scene with different amounts of transformation for a total of 10 images for Gaussian blur and 14 for JPEG compression and uniform light change transformations. Fig. 2 provides

both a sample of the scenes available and an example of transformation applied to an image.

Several well-established datasets, such as [3], are available for evaluating local feature detectors, however are not suitable for use with the proposed framework due to the relatively small number and lesser variety of scenes included, and the limited range of the amount of transformations applied. For example, UBC dataset [3], which was used in [2], includes images for JPEG compression ratios varying from 60% to 98% only. Among the Oxford datasets [3], Leuven offers images under uniform light changes, however the number of images in that dataset is only six. Although the database employed in [16] for assessing several feature detectors under different light conditions contains a large number of images, the number of scenes are limited to 60. Moreover, these scenes were captured in a highly controlled environment so they are lesser representative of real-world scenario in comparison of the image database used here.

The images included in the database utilized for this work have a resolution of 717×1080 pixels and consist of 539 real-

world scenes. Each transformation is applied in several discrete steps to each of the scenes. The Gaussian blur amount is varied in 10 discrete steps from 0 to 4.5σ ($10 \times 539 = 5390$ images), JPEG compression ratio is increased from 0 to 98% in 14 steps. Similarly, the amount of light is reduced from 100% to 10% ($14 \times 539 = 7546$ images). Thus, the database includes a dataset of 10 or 14 images for each of the 539 scenes for a total of 20482 images. The ground truth homography that relates any two images of the same scene is a 3×3 identity matrix for all three transformations as there is no geometric change involved.

Accordingly with the classification criteria introduced in the Section III-C, 51% of the 539 scenes have been labeled as outdoor, 65% contain mostly human made elements and 51% have been attributed with the simple label. Overall, the database has reasonable balance between the content types introduced by the proposed classification criteria, so that it becomes possible to produce significant bias descriptors for the local feature detectors.

V. RESULTS

The proposed framework has been applied for producing the top and lowest rankings for a set of eleven feature detectors which are representative of a wide variety of different approaches [17] and includes the following: Edge-Based Region (EBR) [18], Harris-Affine (HARAFF), Hessian-Affine (HESAFF) [19], Maximally Stable Extremal Region (MSER) [20], Harris-Laplace (HARLAP), Hessian-Laplace (HESLAP) [10], Intensity-Based Regions (IBR) [21], SALIENT [22], Scale-invariant Feature Operator (SFOP) [23], Speeded-Up Robust Feature (SURF) [24] and SIFT [25].

The first subsection provides details on how the repeatability data have been obtained and the second one is dedicated to the discussion about the ranking traits of each local feature detector.

A. REPEATABILITY DATA

The repeatability data are obtained for each transformation type utilizing the image database discussed in Section IV. This data is collected using the authors' original programs with control parameter values suggested by the respective authors. The feature detector parameters could be varied in order to obtain a similar number of extracted features for each detector. However, this has a negative impact on the repeatability of a detector [10] and is therefore not desirable for such an evaluation.

Utilizing the dataset described in detail in Section IV, 18865 repeatability rates have been computed for each local feature detector with the exception of SIFT, which has been assessed only under JPEG compression. It should be noted that SIFT detects more than 20,000 features for some images in the image database which makes it very time-consuming to do such a detailed analysis for SIFT. In the case of JPEG image database, it took more than two months to obtain results on HP ProLiant DL380 G7 system with Intel Xeon 5600 series

processors. Therefore, results for SIFT are not provided in this section.

The number of datasets is 539, the number of discrete step of transformation amount k varies across the transformations considered. We employed: $k = 14$ for JPEG compression and uniform light change transformations and $k = 10$ for Gaussian blur. Since the first step of transformation amount corresponds to the reference image, the number of set B_{kd} (4) is 13 for JPEG compression and light changes, and 9 for blurring for a total of $2 \times (13 \times 539) + 9 \times 539 = 18865$ repeatability rate values for each detector.

B. TRAIT INDICES

In this section the trait indices for all the assessed image feature detectors are presented and discussed. The trait indices have been designed to provide a measure of the bias of the feature detector for any of the types of scene introduced by the classification criterion described in the Section III-C. In other words, they are indicative of the types of scene for which a feature detector is expected to perform well. Accordingly, with the definition provided in the Section III-D, they represent the percentage of the scenes in the top and lowest rankings of a particular type of scene. Thus, they permit to characterize quantitatively the performance of feature detectors from the point of view of the scene content. The trait indices are built starting from the top and lowest rankings of any feature detector. For obtaining the results presented in this work, the evaluation framework has been applied utilizing a ranking length of 20 ($j = 20$). Finally, the related trait indices are computed by applying the equations (7) and (8) presented in the Section III-D.

The results of all detectors are shown in the figures 3–13 and discussed below. The results are presented utilizing radar charts: the transformation amounts are shown on the external perimeter and increase clockwise; the trait indices are expressed in percentage with the value which increases from the centre (%0) to the external perimeter (%100) of the chart.

1) EBR trait indices

All the available trait indices of Edge-Based Regions (EBR) detector [18] are reported in Fig. 3. The performances of EBR are very sensitive to uniform light reduction and Gaussian blur [11]. In particular, for light reduction higher than 60% and Gaussian blur equal or greater than 3.0σ , there are more than 20 scenes for which EBR scores a repeatability rate value of 0% making impossible to form a scene ranking as described above (Section V-B1). For this reason, in Fig. 3.d and 3.e the trait indices for the higher amounts of those transformation are omitted.

EBR exhibits high values (around 80% - 90%) of G in the top rankings and low (rarely above 25%) in the lowest rankings, denoting a strong bias towards the scenes including many human-made elements. EBR performs generally well on simple scenes as well, in particular under Gaussian blur,

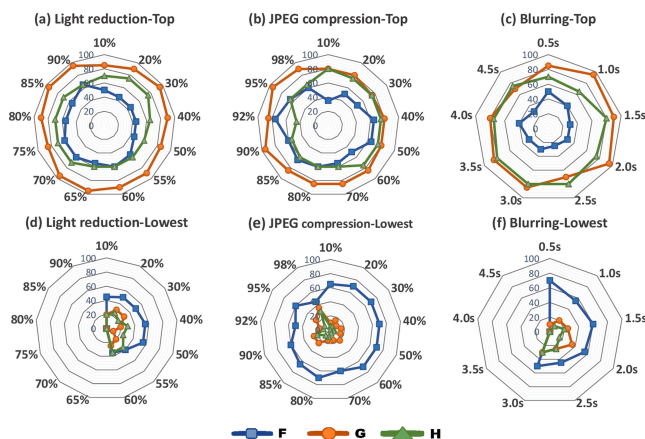


FIGURE 3: Top and lowest trait indices of EBR in percentage for different amounts of light reduction (a,d), JPEG compression (b,e) and blurring (c,f).

with the share of simple scenes never below 70%. The values assumed by F indices are not indicative of the EBR's bias for particular location types as they assume very similar values between the top and lowest rankings.

2) HARLAP and HARAFF trait indices

The rankings of HARLAP [10] and HARAFF [19] are very similar to each other and so are the values of their trait indices. Both of them are particularly prone to uniform light changes and the trait indices for high level of transformation amounts are not available. Fig. 4.a and 5.a report the results for the top rankings up to 80% of light reduction and up to 60% for the lowest rankings.

HARLAP and HARAFF present a bias towards simple scenes, which is particularly strong under uniform light reduction and blurring as can be inferred by the high values, which H assumes in the top twenties and the low values in the relative lowest rankings. A clear preference of those detectors for human made objects can be claimed under light changes

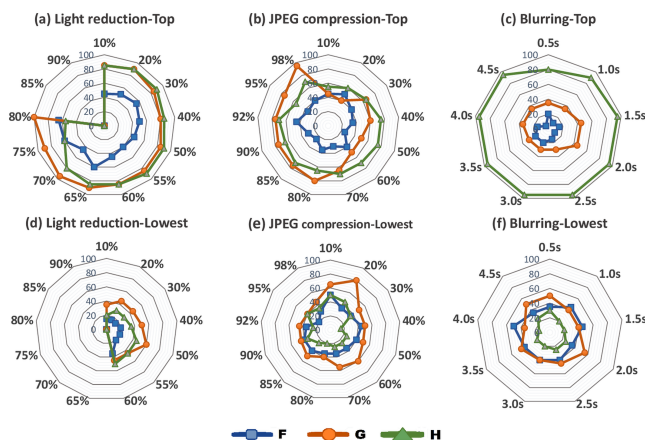


FIGURE 4: Top and lowest trait indices of HARLAP in percentage for different amounts of light reduction (a,d), JPEG compression (b,e) and blurring (c,f).

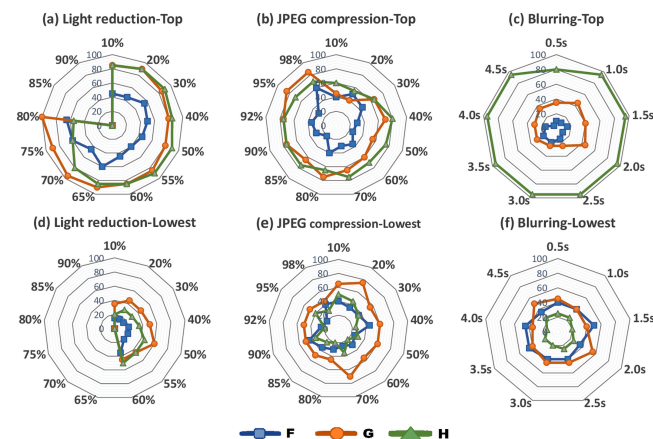


FIGURE 5: Top and lowest trait indices of HARAFF in percentage for different amounts of light reduction (a,d), JPEG compression (b,e) and blurring (c,f).

however, this is not the case under JPEG compression and Gaussian blur whose related G indices are too close between top and lowest rankings to draw any conclusion. The F indices are extremely low (never above 20%) for the top twenty rankings under Gaussian blur revealing that HARLAP and HARAFF deal better with non-outdoor scenes under this particular transformation.

3) HESLAP and HESAFF trait indices

Due to the similarities between the approach in localizing the interest point in images, HESLAP and HESAFF present many similarities between their trait indices. Similarly, to HARLAP and HARAFF, uniform light changes have a strong impact on the HESLAP and HESAFF's performance [11]. For that reason, Fig. 6.a and 7.a show only the results for the top rankings of up to 80% of light reduction and up to 60% for the lowest rankings.

HESLAP and HESAFF perform better on scenes characterized by simple elements and edges under blurring (especially

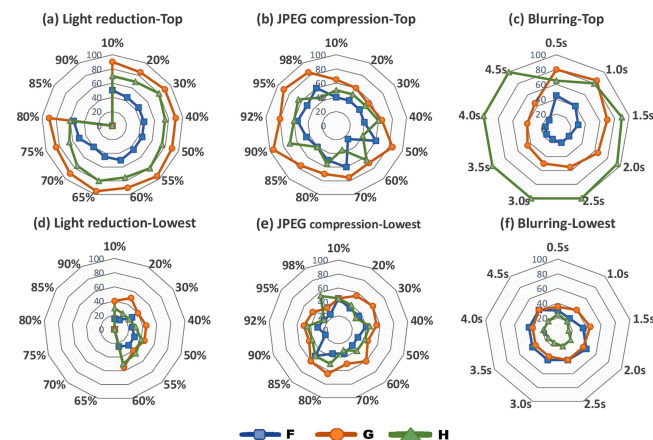


FIGURE 6: Top and lowest trait indices of HESLAP in percentage for different amounts of light reduction (a,d), JPEG compression (b,e) and blurring (c,f).

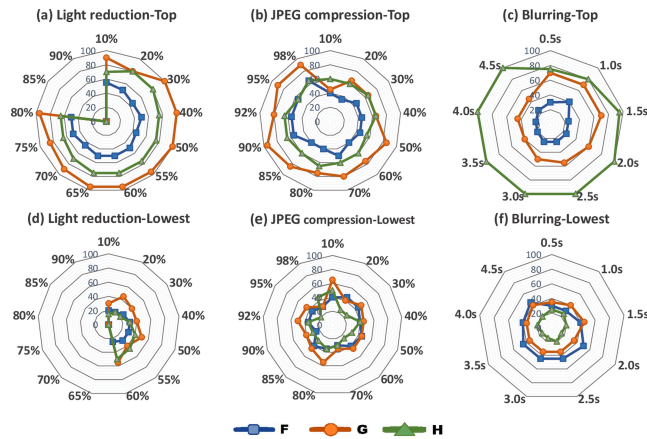


FIGURE 7: Top and lowest trait indices of HESAFF in percentage for different amounts of light reduction (a,d), JPEG compression (b,e) and blurring (c,f).

for high σ values) and uniform light decreasing. The same indices, H , computed under JPEG compression present fluctuations around 50% for both the top and lowest rankings without bending towards simple or complex scenes. Both the detectors perform well on scenes containing human-made elements under light reduction, JPEG compression and up to 2.5σ of Gaussian blur. Although both HESLAP and HESAFF do not have any bias for outdoor scenes, the HASLAP's F index decrease from 45% to 15% constantly over the variation range of blurring amount.

4) SIFT

From the trait index data, it is not possible to determine a clear bias in the performance of SIFT [25], as the values of the trait indices fluctuate over the entire range of the JPEG compression rate. Fig. 8 confirms a bias towards simple and human made objects only between 10% and 60% and at 98% of JPEG compression. whereas the indices G and H present large fluctuations in the top twenty scene rankings for the other compression rates. In particular, between 70% and 90% of JPEG compression their values are significantly lower than ones at other compression amounts and reach a minimum at

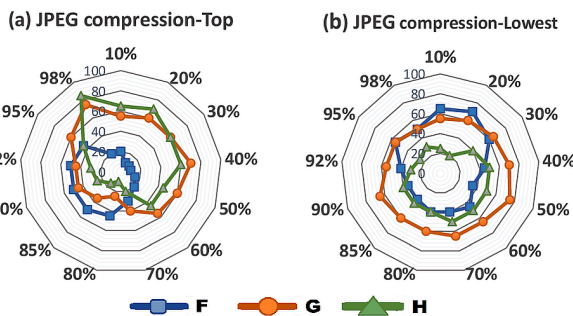


FIGURE 8: Top and lowest trait indices of SIFT in percentage for different amounts JPEG compression (a,b).

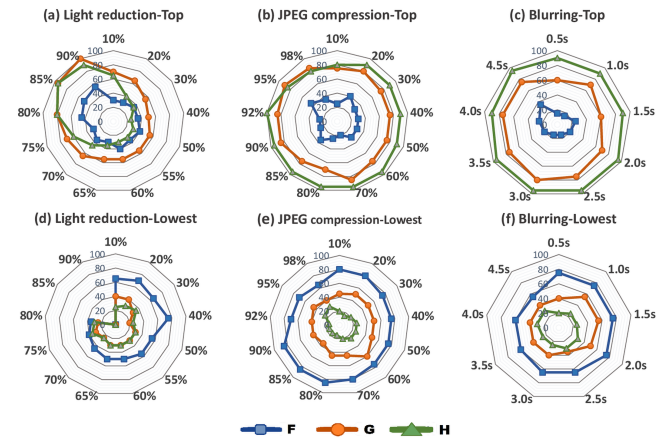


FIGURE 9: Top and lowest trait indices of IBR in percentage for different amounts of light reduction (a,d), JPEG compression (b,e) and blurring (c,f).

80% which are 10% for H and 25% for G . Similar variations can be appreciated also for F in both top and lowest rankings with values variations broad up to 40%. While the G and H indices in many cases present small differences between the top and lowest rankings, the F indices are often inversely related. For example, at 30% compression, F is equal to 10% for the top twenty and 60% for the bottom twenty, G is 60% in both cases and H differs for just 20% between the top and lowest rankings.

In conclusion, the classification criteria adopted in this work permits to infer a strong dependency of SIFT on the JPEG compression rate variations, however, it does not allow to draw any conclusions about the general bias, if any exists, towards a particular type of scene.

5) IBR

The uniform light change has a significant impact on the performance of IBR [21] [11]. This made impossible to obtain the lowest trait indices for brightness reduction at 85% and 90%. Following the same approach as Section V-B1, those indices are set to 0 (Fig. 9.a). Under light reduction, the presence of a weak bias across the range of transformation amount is evident for human-made objects: G indices are never below 50% in top rankings while their counterpart in the lowest indices are never above 40%. A similar trend can be observed for F : the share of outdoor scenes in the top twenty is generally below 50%, while is generally never below 50% for light reduction rates from 10% to 65%. Under JPEG compression, IBR achieved better performances on scenes, which are both simple and human made. Indeed, the related G and H indices in the top twenties reached very high values, which are never below 75% and 80% respectively (Fig. 9.b). The same kind of bias observed for JPEG compression characterized IBR under blurring as well: the top rankings are mainly populated by human made and simple scenes, whereas the lowest rankings contain mostly scenes with the opposite characteristics (Fig. 9.c and 9.f).

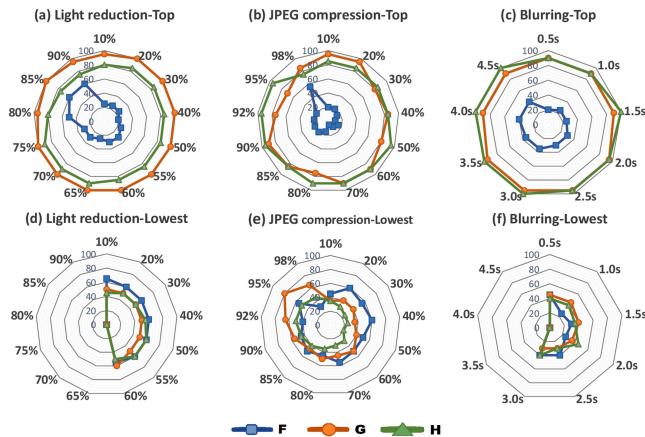


FIGURE 10: Top and lowest trait indices of MSER in percentage for different amounts of light reduction (a,d), JPEG compression (b,e) and blurring (c,f).

6) MSER

Fig. 10 shows the trait indices for MSER [20]. Due to sensitivity of MSER to uniform light reduction and Gaussian blur, it has not been possible to compute the trait indices for the lowest rankings at light reductions of more than 60% and for the last three steps of blurring as the number of scenes with repeatability equal to 0 exceed the length of the lowest rankings at those transformation amounts.

The trait indices draw a very clear picture of the MSER's biases. The very high values of G and H of the top ranking indices and the relatively low values obtained for the lowest twenty rankings, lead to the conclusion that MSER performs significantly better on simple and human-made dominated scenes for every transformation type and amount. Finally, the outdoor scenes populate mainly the lowest rankings built under light reduction and JPEG compression transformations while F for blurring has low and balanced values between the top and lowest rankings.

7) SALIENT

The results for uniform light reduction (Fig. 11.a) shows a strong preference of SALIENT [22] for complex scenes as can be inferred by the low values of the index H in the top twenties contrary to high values in the lowest rankings. This can be explained considering that uniform light reduction does not alter the shape of the edges and other lines present in a scene but makes low contrast region harder to be identified. Thus, the application of the uniform light transformation has the effect of populating the top ranking of SALIENT with those images containing a few but high contrast elements.

On the other hand, the result for Gaussian blur (Fig. 11.c) shows a completely opposite situation, in which the most frequent scenes in the top rankings are those characterized by simple structure or, in other words, scenes whose information has relevant components at low frequencies. Indeed, as

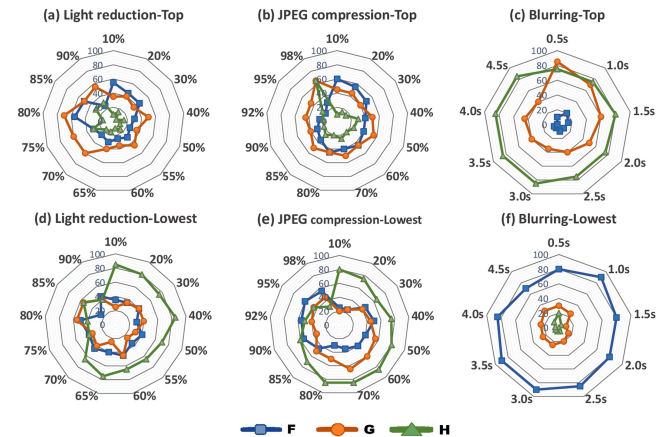


FIGURE 11: Top and lowest trait indices of SALIENT in percentage for different amounts of light reduction (a,d), JPEG compression (b,e) and blurring (c,f).

indicated above, Gaussian blurring acts as a low pass filter and applying it to an image results in loss of the information at high frequencies. Under JPEG compression, SALIENT exhibits a preference for complex scene as it is under light reduction with the difference that the H indices increase with the compression rate. Although, JPEG compression is lossy and may alter the shape of the edges delimiting the potential salient region in an image, the impact on the information content is lower than the one caused by Gaussian blur. Indeed, the share of simple images is constantly low: H below 30% up to 98%. At 98% the share of simple images in the top twenty increases dramatically to 65% as the images lose a huge part of their information content due to the compression, which produce wide uniform regions (see the Fig. 1 for an example).

8) SFOP

Under JPEG compression and Gaussian blur, the bias of SFOP [23] is towards simple scenes representing non-outdoor scenes. The kind of objects favoured are human made under JPEG compression, while for blurring no clear preference can be inferred, due to closeness of the values of G indices between the top and lowest rankings. The relative measures of those biases are reflected by the G and H indices reported in the Fig. 12.b and 12.c: H assumes high percentage values in the top rankings and low values in the lowest rankings; the indices G for the top rankings of JPEG compression are constantly above 70% whereas the related value registered for the lowest rankings exceeds 55% only at 10% of compression rate. The indices obtained for uniform light reduction reveals that SFOP performs lowest on outdoor scenes, as can be seen by the lowest ranking F values, which mostly fluctuate between 50% and 60%.

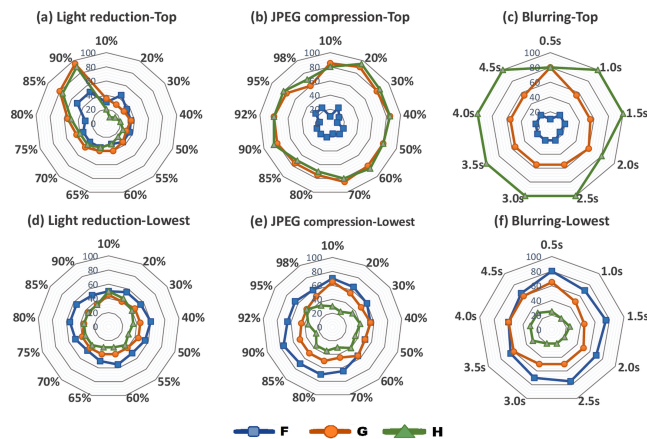


FIGURE 12: Top and lowest trait indices of SFOP in percentage for different amounts of light reduction (a,d), JPEG compression (b,e) and blurring (c,f).

9) SURF

The performance of SURF is particularly affected by uniform light transformation and, because of that, it has not been possible to compute the trait indices at 65% and further brightness reductions (Fig. 13.a). The effect of this transformation is to focus the biases of SURF towards human made objects (G greater than or equal to 65%). Although the available F indices for lowest rankings are extremely low (normally within 15%), only a weak bias towards outdoor scenes can be claimed as the highest values for F indices in the top twenties are only 60% between 10% and 60% of light reduction. The percentage of simple scenes in the top rankings fluctuate between 50% and 85% which, unfortunately, is reached in a region where the indices are not available thus, a comparison is not possible.

JPEG compression produces more predictable biases on SURF: H 's values are significantly higher in the top rankings than in the lowest rankings and the performance are worse with outdoor scenes than with non-outdoor scenes. Finally, the G indices do not express a true bias, neither for human-made nor for natural elements. (Fig. 13.d).

Under blurring SURF best performs on simple scenes whereas it performs poorer on complex scenes. The values of F and G for the top and lowest rankings are fairly close. F values for both rankings groups are very low (except for 0.5σ which reaches 60% for the lowest ranking) while G 's values fluctuate around 50%.

VI. CONCLUSIONS

For several state-of-the-art feature detectors, the dependency of the repeatability from input scene type has been investigated utilizing a large database composed of images from a wide variety of human-classified scenes under three different types and amounts of image transformations. The trait indices collected for the three classes employed by the method allow to infer that the feature detectors tend to score their highest and lower repeatability with particular types of scenes. The detector preferences for a particular category of scene are

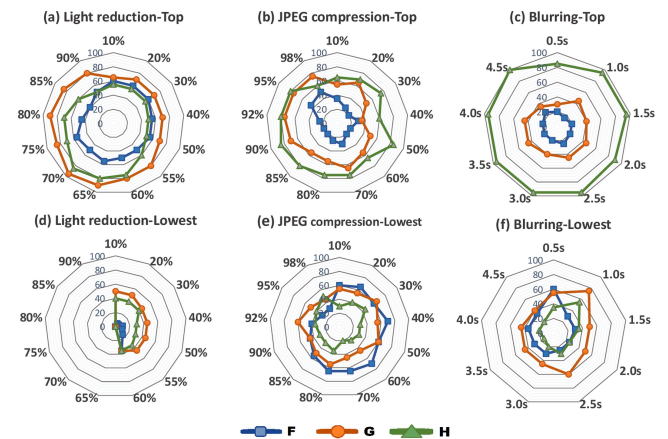


FIGURE 13: Top and lowest trait indices of SURF in percentage for different amounts of light reduction (a,d), JPEG compression (b,e) and blurring (c,f).

pronounced and stable across the type and amount of image transformation for some detectors, such as MSER and EBR. Some detectors' biases are influenced more than others by the amount of transformation and the top-trait indices of SFOP under light changes are a good example: G and H reach a peak at 98% of light reduction. In a few cases the indices show very similar values between top end lower rankings. This allows to conclude that biases of a detector are not sensitive to a particular image change. For example, the trait indices of SALIENT for JPEG compression are between 40% and 60% for most of JPEG compression rate.

A significant number of local image feature detector have been assessed in this work, however the proposed framework is general and can be utilized for assessing any arbitrary set of detectors. A designer who needs to maximize the performance of a vision system starting from the choice of the better possible local feature detector could take advantage from the proposed framework. Indeed, the framework could be utilized for identifying the detectors which perform better with the type of scene most common in the application before any task-oriented evaluation (e.g. [26], [27]) thus, such a selection would be carried out on a smaller set of local feature detectors. For example, for an application which deals mainly with indoor scenes, the detectors should be short-listed are HESAFF, HESLAP, HARAFF and HASAFF which have been proven to achieve their highest repeatability rate with non-outdoor scenes. On the other hand, if an application is intended for working in an outdoor environment, EBR should be one of the considered local feature detectors, especially under light reduction transformation.

In brief, the proposed framework allows to characterize the feature detector against the scene content and, at the same time, represents a useful tool for facilitating the design of those visual applications which utilize a local feature detector stage.

Finally, the work presented in this paper can be extended in several ways. The scene content can be divided in more categories and sub-categories. For example, categories taking

into account 3D structures such as building facades could be added; also, the pattern of the textured surfaces could be described with specific ones. The scene attributes could be represented with fuzzy-like variables instead of binary variables to better reflect intermediate contents. Finally, the behavior of detectors with different scenes could be examined jointly with additional types of image transformation such as point of view and scale changes.

ACKNOWLEDGMENT

This work has been supported by the UK Engineering and Physical Sciences Research Council EPSRC [EP/K004638/1, EP/R02572X/1 and EP/P017487/1]

REFERENCES

- [1] P. Tissainayagam and D. Suter, "Assessing the performance of corner detectors for point feature tracking applications," *Image and Vision Computing*, vol. 22, no. 8, pp. 663–679, 2004.
- [2] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool, "A comparison of affine region detectors," *International Journal of Computer Vision*, vol. 65, no. 1-2, pp. 43–72, 2005.
- [3] K. Mikolajczyk, "Oxford Data Set, <http://www.robots.ox.ac.uk/~vgg/research/affine/>."
- [4] F. Fraundorfer and H. Bischof, "A novel performance evaluation method of local detectors on non-planar scenes," in *Computer Vision and Pattern Recognition-Workshops, 2005. CVPR Workshops. IEEE Computer Society Conference on. IEEE*, 2005, pp. 33–33.
- [5] S. Ehsan, N. Kanwal, A. Clark, and K. McDonald-Maier, "Improved repeatability measures for evaluating performance of feature detectors," *Electronics Letters*, vol. 46, no. 14, pp. 998–1000, 2010.
- [6] S. Ehsan, A. F. Clark, B. Ferrarini, and K. McDonald-Maier, "JPEG, Blur and Uniform Light Changes Image Database, <http://vase.essex.ac.uk/datasets/index.html>."
- [7] B. Ferrarini, S. Ehsan, N. U. Rehman, and K. D. McDonald-Maier, "Performance characterization of image feature detectors in relation to the scene content utilizing a large image database," in *Systems, Signals and Image Processing (IWSSIP), 2015 International Conference on. IEEE*, 2015, pp. 117–120.
- [8] B. Ferrarini, "An Approach to Automatic Selection of the Optimal Local Feature Detector," MSc dissertation, University of Essex, 2016.
- [9] C. Schmid, R. Mohr, and C. Bauckhage, "Evaluation of interest point detectors," *International Journal of Computer Vision*, vol. 37, no. 2, pp. 151–172, 2000.
- [10] K. Mikolajczyk and C. Schmid, "Scale & affine invariant interest point detectors," *International Journal of Computer Vision*, vol. 60, no. 1, pp. 63–86, 2004.
- [11] B. Ferrarini, S. Ehsan, N. U. Rehman, and K. D. McDonald-Maier, "Performance comparison of image feature detectors utilizing a large number of scenes," *Journal of Electronic Imaging*, vol. 25, no. 1, pp. 010 501–010 501, 2016.
- [12] S. Ehsan, A. F. Clark, A. Leonardi, A. Khaliq, M. Fasli, K. D. McDonald-Maier et al., "A generic framework for assessing the performance bounds of image feature detectors," *Remote Sensing*, vol. 8, no. 11, p. 928, 2016.
- [13] F. Fraundorfer and H. Bischof, "Evaluation of local detectors on non-planar scenes," in *In Proc. 28th workshop of the Austrian Association for Pattern Recognition*, 2004.
- [14] T. Dickscheid, F. Schindler, and W. Forstner, "Coding images with local features," *International Journal of Computer Vision*, vol. 94, no. 2, pp. 154–174, 2011.
- [15] S. Ehsan, A. F. Clark, and K. D. McDonald-Maier, "Rapid online analysis of local feature detectors and their complementarity," *Sensors*, vol. 13, no. 8, pp. 10 876–10 907, 2013.
- [16] H. Aanaes, A. L. Dahl, and K. S. Pedersen, "Interesting interest points," *International Journal of Computer Vision*, vol. 97, no. 1, pp. 18–35, 2012.
- [17] T. Tuytelaars and K. Mikolajczyk, "Local invariant feature detectors: a survey," *Foundations and Trends in Computer Graphics and Vision*, vol. 3, no. 3, pp. 177–280, 2008.
- [18] T. Tuytelaars and L. Van Gool, "Content-based image retrieval based on local affinity invariant regions," in *Visual Information and Information Systems*, 1999.
- [19] K. Mikolajczyk and C. Schmid, "An affine invariant interest point detector," in *Computer Vision ECCV 2002. Springer*, 2002, pp. 128–142.
- [20] J. Matas, O. Chum, M. Urban, and T. Pajdla, "Robust wide-baseline stereo from maximally stable extremal regions," *Image and Vision Computing*, vol. 22, no. 10, pp. 761–767, 2004.
- [21] T. Tuytelaars and L. Van Gool, "Matching widely separated views based on affine invariant regions," *International Journal of Computer Vision*, vol. 59, no. 1, pp. 61–85, 2004.
- [22] T. Kadir, A. Zisserman, and M. Brady, "An affine invariant salient region detector," in *ECCV, 2004*, pp. 228–241.
- [23] W. Förstner, T. Dickscheid, and F. Schindler, "Detecting interpretable and accurate scale-invariant keypoints," in *IEEE International Conference on Computer Vision*, 2009.
- [24] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-up robust features (SURF)," *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346–359, 2008.
- [25] D. G. Lowe, "Object recognition from local scale-invariant features," in *IEEE international conference on Computer Vision*, 1999.
- [26] M. C. Shin, D. Goldgof, and K. W. Bowyer, "An objective comparison methodology of edge detection algorithms using a structure from motion task," in *Computer Vision and Pattern Recognition*, 1998. Proceedings. 1998 IEEE Computer Society Conference on. IEEE, 1998, pp. 190–195.
- [27] M. C. Shin, D. Goldgof, and K. W. Bowyer, "Comparison of edge detectors using an object recognition task," in *cvpr. IEEE*, 1999, p. 1360.



BRUNO FERRARINI is a PhD student at the University of Essex, Colchester, UK, where he obtained his MSc (by Dissertation) in Computing and Electronic Systems in 2016. He has an extensive industrial experience in the areas of ICT and InfoSec. His current research interests are in machine learning, local feature detection and description techniques.



SHOAB EHSAN received his BSc Electrical Engineering degree from the University of Engineering and Technology, Taxila, Pakistan in 2003. He completed his PhD in Computing and Electronic Systems (with specialization in computer vision) at the University of Essex, Colchester, UK in 2012. He has extensive industrial and academic experience in the areas of embedded systems, embedded software design, computer vision and image processing. His current research interests are in intrusion detection for embedded systems, local feature detection and description techniques, image feature matching and performance analysis of vision systems. He is a recipient of the University of Essex Post Graduate Research Scholarship and the Overseas Research Student Scholarship. He is a winner of the prestigious Sullivan Doctoral Thesis Prize awarded annually by the British Machine Vision Association.



ALEŠ LEONARDIS is currently Chair of Robotics with the University of Birmingham. He is also an Adjunct Professor with the Faculty of Computer Science, Graz University of Technology. His research interests include robust and adaptive methods for computer vision, object and scene recognition and categorization, statistical visual learning, 3-D object modelling, and biologically motivated vision.



NAVEED UR REHMAN (naveed.rehman@comsats.edu.pk) received the B.Eng. degree in electrical engineering from the National University of Sciences and Technology, Pakistan. He completed his Ph.D. degree in signal processing in 2011 from Imperial College, London, United Kingdom. He is currently working as an Assistant Professor in COMSATS Institute of Information Technology, Islamabad, Pakistan. His research interests include developing data adaptive multivariate T-F algorithms and their applications in biomedical engineering, communications, remote sensing, data fusion, and renewable energy.



PROFESSOR KLAUS D. MCDONALD-MAIER heads the Embedded and Intelligent Systems Laboratory at the University of Essex, Colchester, United Kingdom. He is also Chief Scientist of UltraSoC Technologies Ltd, CEO of Metrarc Ltd and a visiting professor at the University of Kent. His current research interests include embedded systems and system-on-chip (SoC) design, security, development support and technology, parallel and energy efficient architectures, computer vision, data analytics and the application of soft computing and image processing techniques for real world problems. He is a member of the VDE, a senior member of the IEEE and a Fellow of the IET.

...