

In the Saddle: Chasing Fast and Repeatable Features

Javier Aldana-Iuit, Dmytro Mishkin, Ondřej Chum and Jiří Matas
Center for Machine Perception, Czech Technical University in Prague
Prague, Czech Republic.

Email: {aldanjav, mishkdmy, chum, matas}@cmp.felk.cvut.cz

Abstract—A novel similarity-covariant feature detector that extracts points whose neighborhoods, when treated as a 3D intensity surface, have a saddle-like intensity profile. The saddle condition is verified efficiently by intensity comparisons on two concentric rings that must have exactly two dark-to-bright and two bright-to-dark transitions satisfying certain geometric constraints.

Experiments show that the Saddle features are general, evenly spread and appearing in high density in a range of images. The Saddle detector is among the fastest proposed. In comparison with detector with similar speed, the Saddle features show superior matching performance on number of challenging datasets.

I. INTRODUCTION

Local invariant features¹ have a wide range of applications: image alignment and retrieval [1], [2], specific object recognition [3], [4], 3D reconstruction [5], [6], robot location [7], tracking by detection [8], augmented reality [9], etc. It is therefore not surprising that literally hundreds of local feature detectors have been proposed [10], [11].

In an application, the suitability of a particular local feature detector depends typically on more than one property. The important characteristics are most commonly the *repeatability* – the ability to respond to the same scene pre-image irrespective of changing acquisition conditions, *distinctiveness* – the discriminative power of the intensity patches it extracts, *density* – the number of responses per unit area, both average and maximum achievable and *efficiency* – the speed with which the features are extracted. Other properties like the *generality* of the scenes where the feature exhibits acceptable performance of the major characteristics, the evenness of the *coverage* of image, the geometric *accuracy* are considered less often.

Local feature detectors with the largest impact lie on the "convex envelope" of the properties. The Difference-of-Gaussians [3] and the Hessian, either in the rotation [12], similarity [10] or affine covariant [13] form, are arguably the most general detectors with high repeatability [14]. For their efficiency, SURF [15], FAST [16] and ORB [17] are the preferred choice for real-time applications or in cases when computational resources are limited as on mobile devices. MSERs [18] are popular for matching of images with extreme viewpoint changes [19] and in some niches like text detection [20], [21]. Learned detectors, trained to specific requirements like insensitivity to gross illumination changes, outperform, in their domains, generic detectors [22]. For some problems, like matching between different modalities, any

¹a.k.a. interest points, keypoints, feature points, distinguished regions.

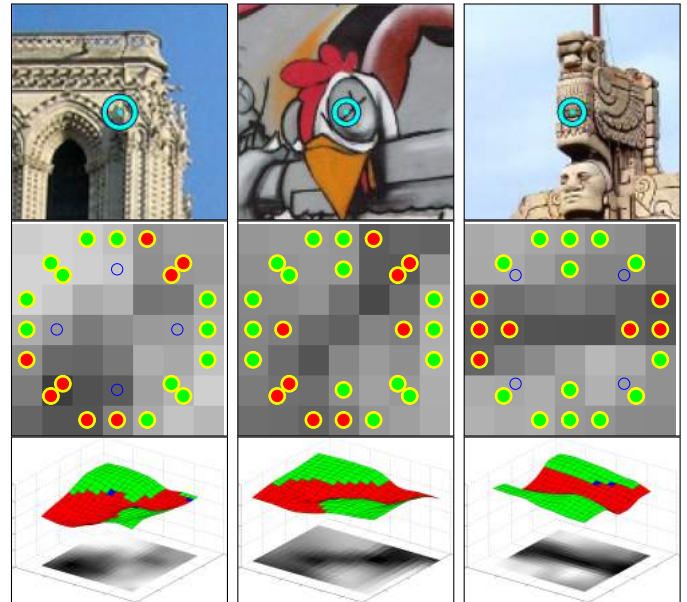


Fig. 1. Saddle feature examples (top row). Corresponding image patches with accepted arrangements of dark (marked red), bright (green) and intermediate (blue) pixel intensities (middle row). Pixel intensities around Saddle points visualized as a 3D surface (bottom row).

single detector is inferior to a combination of different local feature detectors [23].

As a necessary condition of an interest point [24], the patch around the interest point must be dissimilar to patches in its immediate neighborhood. There are at least three types of such interest regions: (i) corners such as Harris corner detector [24], (ii) blobs such as MSER [18], DoG [3] or Hessian with positive determinant [10], and (iii) saddle points, e.g. Hessian with negative determinant [10]. Rapid detectors of corner points FAST [16] and ORB [17] and of blobs SURF [15] have been already proposed and are used in applications with significant time constraints.

In this paper, we propose a novel similarity-covariant local feature detector called Saddle. The detector extracts points whose neighborhoods, when treated as a 3D intensity surface, have concave and convex profiles in a pair of orthogonal directions, see Fig. 1; in a continuous setting the points would have a negative determinant of the Hessian matrix. The saddle condition is verified on two concentric approximately circular rings which must have exactly two dark-to-bright and two bright-to-dark transitions satisfying certain geometric constraints, see Fig. 2.




Fig. 2. The 8 pixel positions marked red form the *inner ring* and the 16 positions b_j marked blue form the *outer ring*. Positions shared by both rings are bicolored.

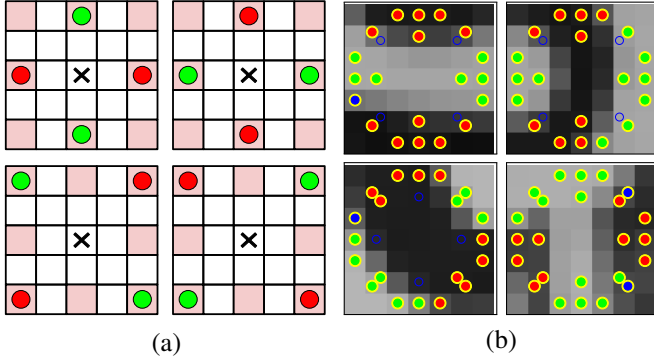


Fig. 3. (a) The fast test for an alternating pattern on the inner ring. In each of the four patterns, green dots depict pixels with intensity strictly brighter than the intensity of pixels marked red. The location is eliminated further by Saddle if none of the patterns is observed. (b) Examples of accepted patterns.

Experiments show that such points exist with high density in a broad class of images, are repeatably detectable, distinctive and are accurately localized. The Saddle points are stable with respect to scale and thus a coarse pyramid is sufficient for their detection, saving time and memory. Saddle is faster than SURF, a popular choice of detector when fast response is required, but slower than ORB. Overall, the Saddle detector provides an attractive combination of properties sufficient to have impact even in the mature area of local feature detectors.

Saddle falls into the class of detectors that are defined in terms of intensity level comparisons, together with BRISK [25], FAST [16], its similarity-covariant extension ORB [17], and its precursors like SUSAN [26] and the Trajkovic-Hedley detector [27]. With the exception of BRISK, the intensity-comparison based detector aim at corner-like features and can be interpreted as a fast approximation of the Harris interest point² detector [24]. Saddle is novel in that it uses intensity comparisons for detection of different local structures, related to Hessian rather than the Harris detector.

II. THE SADDLE INTEREST POINT DETECTOR

The algorithmic structure of the Saddle keypoint detector is simple. Covariance with similarity transformation is achieved by localizing the keypoints in a scale-space pyramid [28]. At every level of the pyramid, the Saddle points are extracted in three steps. First, a fast alternating pattern test is performed in the inner ring, see Figs. 2 and 3. This test eliminates about 80–85% of the candidate points. If a point passes the first

²In fact, the ORB final interest point selection is a function of the Harris response computed on points that pass a preliminary test.

Algorithm 1 Saddle feature detection

Input: Image I , ϵ
Output: Set F of Saddle keypoints

```

for pyramid level  $I_n$  do
  for every pixel  $\mathbf{p}$  in  $I_n$  do
    if  $\sim$  INNER( $\mathbf{p}$ ) then
      continue
    Compute  $\rho_{\mathbf{p}}$ 
    if  $\sim$  OUTER( $\mathbf{p}$ ,  $\rho_{\mathbf{p}}$ ,  $\epsilon$ ) then
      continue
    Compute response  $R(\mathbf{p})$ 
  Non-Maxima Suppression
  Coordinate Refinement
return

```




Fig. 4. Detection on a progressively blurred chessboard pattern. Circle color reflects feature scale, its size shows the extent of the description region.

test, an alternating pattern test on the outer ring is carried out. Finally, points that pass both tests enter the post-processing stage, which includes non-maxima suppression and response strength selection. The algorithm is summarized in Alg. 1.

A. Alternating pattern on the inner ring

The first test is designed to be very fast and to reject majority of points. The test operates on pixels surrounding the central point – the pink square in Fig 2. In the test, two pairs of orthogonal directions are considered, one in the shape of + and the other in the shape of \times . The test is passed if both points on the inner ring in one direction are strictly brighter than both points in the orthogonal direction. The four cases for passing the test are depicted in Fig. 3 (a). Note that either of the + and \times shapes can pass the test, or both.

From the intensity values of the pixels satisfying the inner pattern test, either four or eight pixels, depending whether one or both patterns passed the test, central intensity value ρ is estimated. As a robust estimate, the median of the intensity values is used.

B. Alternating pattern on the outer ring

The second test considers the 16 pixels that approximate a circle of radius 3 around the central point. The outer ring is depicted in Fig. 2 in light blue. Let the pixels on the outer ring

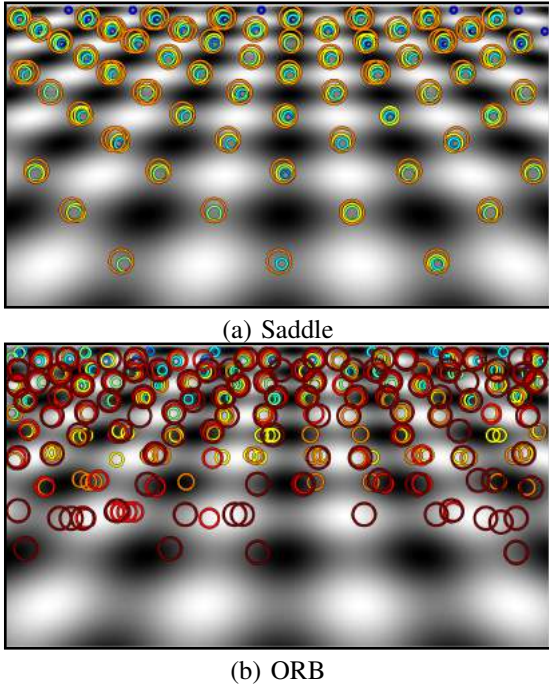


Fig. 5. Detection on a 2D sinusoidal pattern under a perspective transformation. Saddle and ORB detections are shown as circles of the outer ring size.

be denoted as $B = \{b_j \mid j = 1 \dots 16\}$. Each of the pixels in B is labeled by one of three labels $\{d, s, l\}$. The labels are determined by the pixel intensity I_{b_j} , the central intensity at the saddle point ρ , and the method parameter offset ε as follows

$$L_{b_j} = \begin{cases} \bullet d, & I_{b_j} < \rho - \varepsilon \\ \bullet s, & \rho - \varepsilon \leq I_{b_j} \leq \rho + \varepsilon \\ \bullet l, & I_{b_j} > \rho + \varepsilon \end{cases} \quad (1)$$

The color of the dots in (1) corresponds to the color of the dots in the outer ring in Figs. 1 and 3 (b).

The test is passed if the outer ring contains exactly two consecutive arcs of each label l and d , the arcs are of length 2 to 8 pixels and are alternating – the l arcs are separated by d arcs. To eliminate instability caused by ρ -crossing between l and d arcs, up to two pixels can be labeled s at each boundary between l and d arcs. Labels s are pixels with intensity in ε -neighborhood of ρ , where ε is a parameter of the detector.

The test may seem complex, but in fact it is a regular grammar expression, which is equivalent to a finite state automaton and can be implemented very efficiently.

C. Post-processing

Each point \mathbf{p} that passed the alternating pattern test for both the inner and outer ring is assigned a response strength

$$R(\mathbf{p}) = \sum_{b_j \in B(\mathbf{p})} |\rho_{\mathbf{p}} - b_j|.$$

The value of the response strength is used in the non-maxima suppression step and to limit the number of responses if required.

The non-maxima suppression is only performed within one level of the pyramid, features at different scales do not interact as the scale pyramid is relatively coarse. This is similar to non-maxima suppression of ORB. For the non-maxima suppression, a 3×3 neighborhood of point \mathbf{p} is considered.

As a final post-processing step, position refinement of points that passed the non-maxima suppression state takes place. A precise localization of the detected keypoint \mathbf{p} within the pyramid level is estimated with sub-pixel precision. The x and y coordinates of \mathbf{p} are computed as a weighted average of coordinates over a 3×3 neighborhood, where the weights are the response strengths R of each pixel in the neighborhood. Response of pixels that do not pass the alternating pattern tests is set to 0. The feature orientation is defined by the vector from the feature center to the intensity centroid [30], computed within the image patch of 31×31 pixels in the corresponding scale.

III. EXPERIMENTS

In this section, we experimentally evaluate the properties of the proposed Saddle detector. The performance is compared with a number of commonly used feature detectors on standard evaluation benchmarks.

A. Synthetic images

Properties of the Saddle and ORB, first are compared in two experiments on synthetically generated images.

First, features are detected on a chessboard pattern with progressively increasing blur, see Fig. 4. Saddle point detection is expected in the central strips, ORB detection on the corners on the right edge and potentially near the saddle points. Saddle features are repeatedly detected at all blur levels and are well located at the intersection of the pattern edges. ORB features are missing at higher blur levels and their position is less stable.

A phenomenon common to corner feature points – shifting from the corner for higher scales and blur levels is also visible. Note that since the scaling factor between pyramid levels of Saddle is 1.3 while for ORB it is 1.2, Saddle is run on a 6 level pyramid and ORB with 8 to achieve a similar range of scales.

Second, a standard synthetic test image introduced by Lindenberg and used in scale-space literature [28] is used, see Fig. 5. The Saddle points are output at locations corresponding to saddle points across all scales in the perspectively distorted $f(\xi, \eta) = \sin(\xi) \sin(\eta)$ pattern. Since there are no corners in the image, ORB detections are far from regular and are absent near the bottom edge. Fig. 7 shows the detector complementarity, i.e. Saddle fires on regions where other detectors have none detections.

B. Matching coverage

In some task, such as structure from motion, good coverage of the image by matched point is crucial for the stability of the geometric models and consequently for the reliability of the 3D reconstruction [31]. Note that the coverage is a complementary criterion to the number of matched features, which is

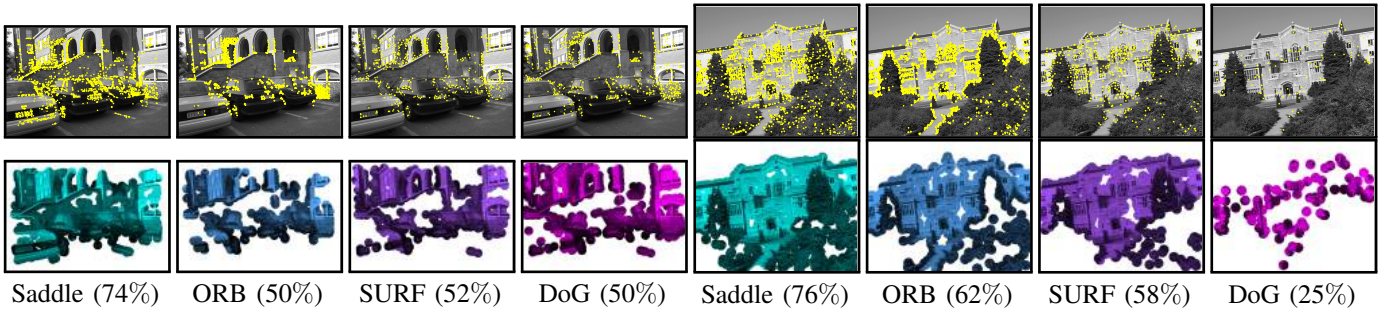


Fig. 6. Coverage by ground-truth validated feature matches on selected image pairs from the Oxford dataset [14], [29]. Yellow dots mark positions of the features (top). The covered area is computed as a union of circles with a 25 pixel radius centered on the matches (bottom).




Fig. 7. Positions of matched interest regions detected with Saddle, ORB, SURF and DoG showing the detection complementarity.

addressed in Section III-D. A high number of clustered matches may lead to poor geometry estimation and to incomplete 3D reconstruction.

To compare the coverage of different feature detectors, we adopt the measure proposed in [31]. An image coverage mask is generated from matched features. Every tentative correspondence geometrically consistent with the ground truth homography adds a disk of a fixed radius (of 25 pixels) into the mask at the location of the feature point. The disk size does not change with the scale of the feature. The matching coverage is then measured as a fraction of the image covered by the coverage mask.

Extensive experiments show that the proposed Saddle detector outperforms all other compared detectors: ORB, SURF and DoG. Quantitative results are shown in Figure 8. The covered areas are shown in Figure 6. The superior coverage of the Saddle detector is visible on Fig. 11.

C. Accuracy

The accuracy of Saddle was assessed on the Oxford-Affine dataset. The cumulative distributions of reprojection errors with respect to the ground truth homography are shown in Figure 9. Saddle marginally outperforms ORB and DoG performance is superior in most cases.

D. Matching ability

In this section we follow the detector evaluation protocol from [23]. We apply it to a restricted number of detectors – those that are direct competitors of Saddle: ORB [17], Hessian [10] (extracting similar keypoints) and SURF [32] (also known as FastHessian).

We focus on getting a reliable answer to the match/no-match question for challenging image pairs. Performance is therefore measured by the number of successfully matched pairs, i.e.

TABLE I
DATASETS USED IN EVALUATION

Short name	Proposed by	#images	Nuisance type
OxAff	Mikolajczyk <i>et al.</i> [14], [29], 2013	8x6	Geom., blur, illum.
EF	Zitnick and Ramnath <i>et al.</i> [33], 2011	8x6	geom., blur, illum.
GDB	Kelman <i>et al.</i> [34], 2007	22x2	illum., sensor
Symb	Hauage and Snavely [35], 2012	46x2	appearance

those with at least 15 inliers found. The average number of inliers provides a finer indicator of the performance.

The datasets used in this experiment are listed in Table I. Results are presented in two tables. Table II shows the results for a setup that focuses on matching speed and thus uses the fast BRIEF [36] and FREAK [37] descriptors (OpenCV implementation). Saddle works better with FREAK, while ORB results are much better with BRIEF. Saddle covers larger area and on broad class of images (e.g. see Figure 11), but needs different descriptor than BRIEF, possible optimized for description of saddle points.

In an experiment Saddle is run with a combination of Root-SIFT [38] and HalfRootSIFT [34] as descriptors (see Table II). This combination was claimed in recent benchmark [23] as best performing along broad range of datasets and it is suitable for evaluation of the matching potential of the feature detectors. With the powerful descriptors, Saddle clearly outperforms ORB. The MODS and WxBS are added as state-of art matchers in their original setup. Most time is taken by description and matching.

Note that one could use both Saddle and ORB detectors and benefit both from their speed and their complementarity (last rows in Table II).

E. Speed

The time breakdown for Saddle and ORB image matching on the Oxford-Affine dataset is shown in Fig 10. Saddle is about two times slower than ORB in the detection part. However, we have neither utilized SSE instructions in the Saddle tests. The results show that both Saddle and ORB are faster than the FREAK descriptor, but significantly slower than BRIEF. The slower RANSAC step for ORB with BRIEF is due to the lower inlier ratio. For each of the two cross inner tests of Saddle, 75% of the points require 2 comparisons, 5% require 3 and 20% require 4. 15% of the points are accepted after 4 comparisons. Percentages refer to the total number of neighbourhoods tested.

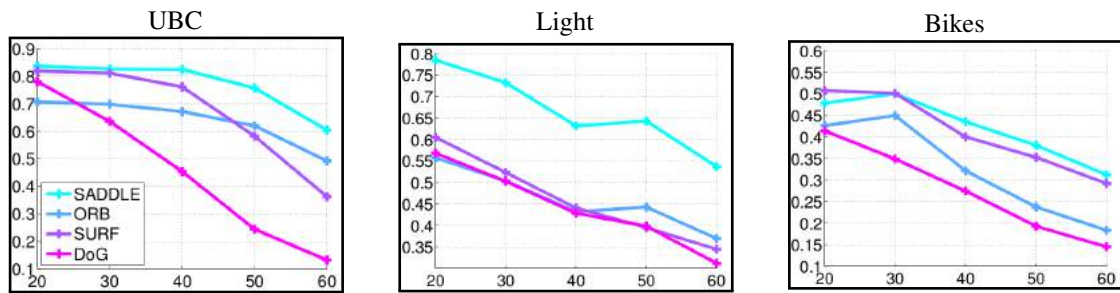


Fig. 8. Coverage by ground-truth validated feature matches on six image sets from the Oxford-Affine dataset [14], [29]. The x-axis shows the viewpoint angle and the y-axis shows the inlier coverage ratio in the reference image.

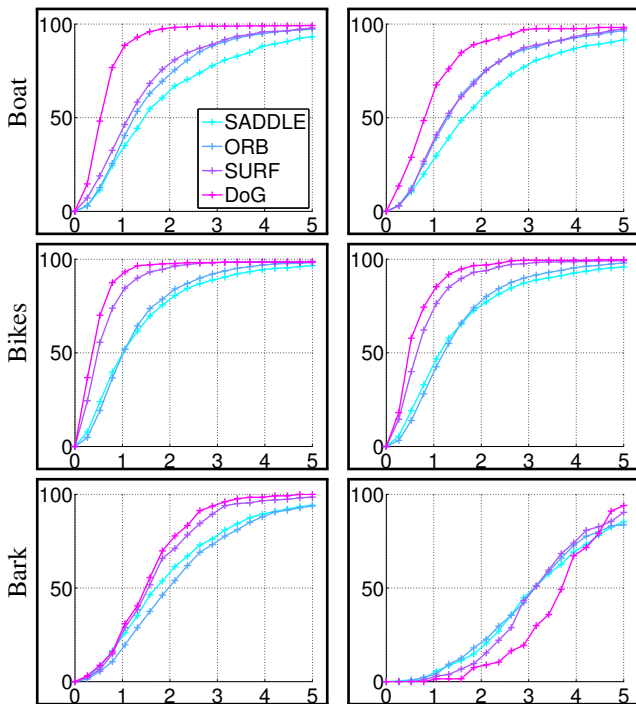


Fig. 9. Inlier ratio (y-axis) curves on Oxford dataset [14]. The reprojection error is given in pixels (x-axis) in the reference image.




Fig. 10. Average run-time for ORB and Saddle on the Oxford-Affine dataset (average image size is $\approx 900 \times 600$ and average number of features is ≈ 1000).

IV. CONCLUSION

In this work we presented Saddle – a novel similarity-covariant feature detector that responds to distinctive image regions at saddle points of the intensity function.

Experiments show that the Saddle features are general, evenly spread and appearing in high density in a range of images. The Saddle detector is among the fastest proposed. In comparison with detector with similar speed, the Saddle features show superior matching performance on number of challenging datasets.

ACKNOWLEDGMENT

Javier Aldana-Iuit was supported by CONCIYTEY-CONACYT-Mexico PhD scholarship 216786 and CTU student grant SGS15/155/OHK3/2T/13, Dmytro Mishkin by CTU student grant SGS15/155/OHK3/2T/13, Ondřej Chum by MSMT LL1303 ERC-CZ grant, and Jiří Matas by The Czech Science Foundation Project GACR P103/12/G084.

REFERENCES

- [1] C. Schmid and R. Mohr, "Local grayvalue invariants for image retrieval," *PAMI*, 1997.
- [2] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman, "Object retrieval with large vocabularies and fast spatial matching," in *CVPR*, 2007.
- [3] D. G. Lowe, "Object recognition from local scale-invariant features," in *ICCV*, 1999.
- [4] S. Obdrzalek and J. Matas, "Sub-linear indexing for large scale object recognition," in *BMVC*, 2005.
- [5] J.-M. Frahm, M. Pollefeys, S. Lazebnik, B. Clipp, D. Gallup, R. Raguram, and C. Wu, "Fast robust reconstruction of large-scale environments," in *CISS*, 2010.
- [6] S. Agarwal, N. Snavely, I. Simon, S. M. Seitz, and R. Szeliski, "Building rome in a day," in *ICCV*, 2009.
- [7] S. Se, D. G. Lowe, and J. J. Little, "Mobile robot localization and mapping with uncertainty using scale-invariant visual landmarks," *IJRR*, 2002.
- [8] M. Özuysal, V. Lepetit, F. Fleuret, and P. Fua, "Feature harvesting for tracking-by-detection," in *ECCV*, 2006.
- [9] V. Lepetit, L. Vacchetti, D. Thalmann, and P. Fua, "Fully automated and stable registration for augmented reality applications," in *ISMAR*, 2003.
- [10] K. Mikolajczyk and C. Schmid, "Scale & affine invariant interest point detectors," *IJCV*, 2004.
- [11] T. Tuytelaars and K. Mikolajczyk, "Local invariant feature detectors: A survey," *FTCVG*, 2008.
- [12] P. R. Beaudet, "Rotationally invariant image operators," in *IJCP*, 1978.
- [13] K. Mikolajczyk and C. Schmid, "An affine invariant interest point detector," in *ECCV*, 2002.
- [14] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool, "A comparison of affine region detectors," *IJCV*, 2005.

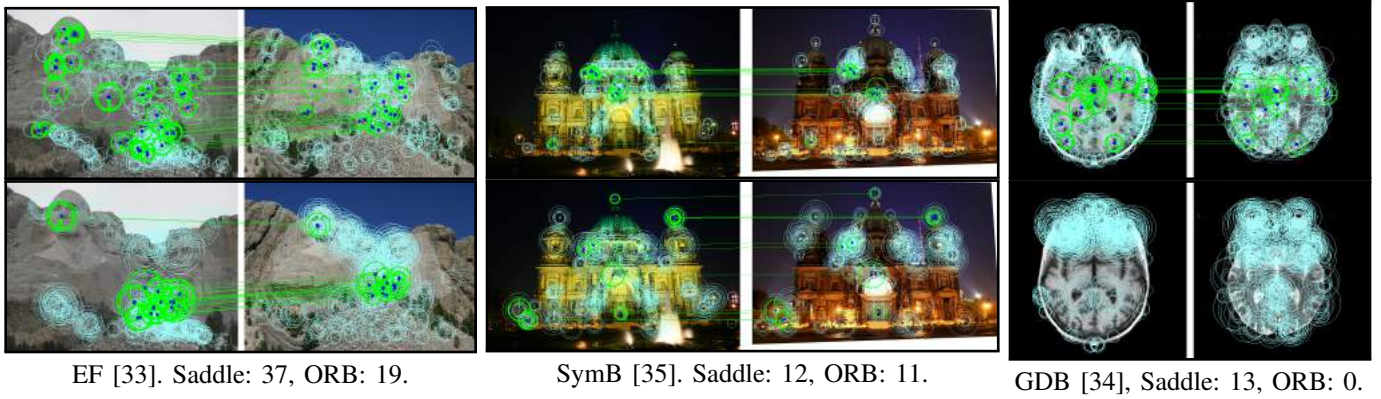


Fig. 11. Detected and matched keypoints for Saddle (top) and ORB (bottom). The inliers count is given for both detectors for each image. Note that Saddle points are spread more evenly making the homography estimation more stable. Images were selected from datasets listed in Table I

TABLE II

SADDLE EVALUATION WITH FAST BRIEF AND FREAK DESCRIPTORS. THE SUBCOLUMNS ARE: THE NUMBER OF SUCCESSFULLY MATCHED IMAGE PAIRS (LEFT), AVERAGE RUNNING TIME (ALL STAGES: READ IMAGE-DETECT-DESCRIBE-MATCH-RANSAC), AVERAGE NUMBER OF INLIERS IN MATCHED PAIRS (RIGHT). B STANDS FOR BRIEF, F – FOR FREAK. THE DATASETS ARE LISTED IN TABLE I. DARKER CELL BACKGROUND INDICATES BETTER RESULTS.

Alg.	Sens	Desc	EF			OxAff			SymB		
			33 #	time [s]	inl. #	40 #	time [s]	inl. #	46 #	time [s]	inl. #
Saddle	0.5K	B	2	0.3	32	22	0.3	47	1	0.3	41
Saddle	0.5K	F	11	0.3	29	26	0.4	77	8	0.4	32
ORB	0.5K	B	16	0.1	28	27	0.1	110	14	0.2	30
ORB	0.5K	F	0	0.1	0	0	0.5	0	0	0.2	0
Saddle	1K	B	2	0.3	52	25	0.3	92	3	0.3	70
Saddle	1K	F	14	0.4	47	29	0.4	145	11	0.4	59
ORB	1K	B	18	0.2	52	28	0.2	208	21	0.2	51
ORB	1K	F	0	0.2	0	0	0.3	0	0	0.2	0
Saddle+ORB	0.5K	B	7	0.3	51	27	0.4	140	11	0.4	41
Saddle+ORB	0.5K	F	19	0.4	42	29	0.6	170	14	0.5	51
Saddle+ORB	1K	B	9	0.3	75	28	0.4	269	12	0.4	82
Saddle+ORB	1K	F	18	0.4	87	31	0.9	317	21	0.5	85

TABLE III

SADDLE EVALUATION WITH A COMBINATION OF ROOTSIFT AND HALFRSIFT DESCRIPTORS. THE SUBCOLUMNS ARE THE SAME AS IN TABLE IV. NMS STANDS FOR SPATIAL NON-MAXIMUM SUPPRESSION, INDICATING ITS APPLICATION. IN MODS-S, ORB WAS REPLACED BY SADDLE+FREAK, OTHER PARAMETERS KEPT ORIGINAL. DARKER CELL BACKGROUND INDICATES BETTER RESULTS.

Alg.	Sens	NMS	Desc	EF			OxAff			SymB		
				33 #	time [s]	inl. #	40 #	time [s]	inl. #	46 #	time [s]	inl. #
Saddle	0.5K	-	SIFT	9	0.5	33	30	0.6	73	8	0.6	31
ORB	0.5K	-	SIFT	13	0.7	35	36	0.8	112	12	0.8	46
Saddle	0.5K	+	SIFT	18	0.5	46	32	0.6	123	15	0.6	49
ORB	0.5K	+	SIFT	12	0.7	37	36	0.9	112	13	0.8	44
Saddle	1K	-	SIFT	18	0.9	49	33	0.9	134	17	0.9	48
ORB	1K	-	SIFT	25	1.5	49	37	1.5	222	27	1.4	55
Saddle	1K	+	SIFT	25	0.9	70	34	0.9	249	25	1.0	71
ORB	1K	+	SIFT	24	1.4	50	37	1.5	222	27	1.5	55
SURF			SIFT	31	0.6	51	37	1.3	483	35	1.2	93
MODS			SIFT	33	1.0	36	40	0.5	163	41	3.9	35
MODS-S			SIFT	33	1.8	34	40	1.7	257	43	4.6	40
WXBS			SIFT	33	1.2	41	40	1.2	149	45	5.5	45

[15] H. Bay, T. Tuytelaars, and L. V. Gool, "Surf: Speeded up robust features," in *ECCV*, 2006.

[16] E. Rosten and T. Drummond, "Machine learning for high-speed corner detection," in *ECCV*, 2006.

[17] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *ICCV*, 2011.

[18] J. Matas, O. Chum, M. Urban, and T. Pajdla, "Robust wide baseline stereo from maximally stable extrema regions," in *BMVC*, 2002.

[19] D. Mishkin, J. Matas, and M. Perdoch, "Mods: Fast and robust method for two-view matching," *CVIU*, 2015.

[20] L. Neumann and J. Matas, "On combining multiple segmentations in scene text recognition," in *ICDAR*, 2013.

[21] Y. Li, W. Jia, C. Shen, and A. van den Hengel, "Characterness: an indicator of text in the wild," *IP, IEEE*, 2014.

[22] Y. Verdie, K. M. Yi, P. Fua, and V. Lepetit, "TILDE: A Temporally Invariant Learned DEtector," in *CVPR*, 2015.

[23] D. Mishkin, J. Matas, M. Perdoch, and K. Lenc, "WxBS: Wide Baseline Stereo Generalizations," in *BMVC*, 2015.

[24] C. Harris and M. Stephens, "A combined corner and edge detector," in *AVC*, 1988.

[25] S. Leutenegger, M. Chli, and R. Y. Siegwart, "Brisk: Binary robust invariant scalable keypoints," in *ICCV*, 2011.

[26] S. M. Smith and J. M. Brady, "Susan: A new approach to low level image processing," *IJCV*, 1997.

[27] M. Trajcovic and M. Hedley, "Fast corner detection," *IVC*, 1998.

[28] T. Lindeberg, "Discrete scale-space theory and the scale-space primal sketch," Ph.D. dissertation, Royal Inst. of Technology, Stockholm, Sweden, 1991.

[29] K. Mikolajczyk and C. Schmid, "A performance evaluation of local descriptors," *PAMI*, 2005.

[30] P. L. Rosin, "Measuring corner properties," *CVIU*, 1999.

[31] A. Irschara, C. Zach, J.-M. Frahm, and H. Bischof, "From structure-from-motion point clouds to fast location recognition," in *CVPR*, 2009.

[32] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-up robust features (surf)," *CVIU*, 2008.

[33] C. L. Zitnick and K. Ramnath, "Edge foci interest points," in *ICCV*, 2011.

[34] A. Kelman, M. Sofka, and C. V. Stewart, "Keypoint descriptors for matching across multiple image modalities and non-linear intensity variations," in *CVPR*, 2007.

[35] D. Hauage and N. Snavely, "Image matching using local symmetry features," in *CVPR*, 2012.

[36] M. Calonder, V. Lepetit, C. Strecha, and P. Fua, "Brief: Binary robust independent elementary features," in *ECCV*, 2010.

[37] A. Alahi, R. Ortiz, and P. Vandergheynst, "FREAK: Fast Retina Keypoint," in *CVPR*, 2012.

[38] R. Arandjelović and A. Zisserman, "Three things everyone should know to improve object retrieval," in *CVPR*, 2012.