



submitted to Pattern Recognition Letters  
journal homepage: [www.elsevier.com](http://www.elsevier.com)

## HarrisZ<sup>+</sup>: Harris Corner Selection for Next-Gen Image Matching Pipelines

Fabio Bellavia<sup>a,\*\*</sup>, Dmytro Mishkin<sup>b</sup>

<sup>a</sup>Dept. of Mathematics and Computer Science, Università Degli Studi di Palermo, Italy

<sup>b</sup>Visual Recognition Group, Faculty of Electrical Engineering, CTU in Prague

### ABSTRACT

Due to its role in many computer vision tasks, image matching has been subjected to an active investigation by researchers, which has led to better and more discriminant feature descriptors and to more robust matching strategies, also thanks to the advent of the deep learning and the increased computational power of the modern hardware. Despite of these achievements, the keypoint extraction process at the base of the image matching pipeline has not seen equivalent progresses. This paper presents HarrisZ<sup>+</sup>, an upgrade to the HarrisZ corner detector, optimized to synergically take advantage of the recent improvements of the other steps of the image matching pipeline. HarrisZ<sup>+</sup> does not only consist of a tuning of the setup parameters, but introduces further refinements to the selection criteria delineated by HarrisZ, so providing more, yet discriminative, keypoints, which are better distributed on the image and with higher localization accuracy. The image matching pipeline including HarrisZ<sup>+</sup>, together with the other modern components, obtained in different recent matching benchmarks state-of-the-art results among the classic image matching pipelines. These results are quite close to those obtained by the more recent fully deep end-to-end trainable approaches and show that there is still a proper margin of improvement that can be granted by the research in classic image matching methods.

### 1. Introduction

Image keypoint extraction has played a relevant role in computer vision since the early days (Szeliski, 2021). A keypoint is generally defined in a broad sense as a local region on the image which can be correctly re-localized and distinguished from others after a transformation of the image, i.e. keypoints must be repeatable and discriminable. The precision and the computational efficiency of the keypoint localization, the sparseness of the keypoint distribution over the images, and the kind of image transformations the keypoint must tolerate generally depend on the application purpose (Tuytelaars and Mikolajczyk, 2008).

With the increasing success of deep learning approaches in replacing the handcrafted ones, keypoint extraction more or less implicitly has been hidden within the lay-

ers of end-to-end trainable networks (Dusmanu et al., 2019; Tian et al., 2020; DeTone et al., 2018; Sun et al., 2021) or apparently removed (Truong et al., 2021). Nevertheless, keypoint detectors as standalone building blocks are still actively investigated and successfully used as the base of modern and competitive image matching pipeline for 3D reconstruction in sparse Structure-from-Motion (SfM, Schönberger and Frahm (2016)) and Simultaneous Localization and Mapping (SLAM, Mur-Artal et al. (2015)).

Although besieged by deep keypoint detectors, handcrafted detectors are still able to provide state-of-the-art results in SfM and SLAM applications (Jin et al., 2020). The keypoint extraction is the first step of 3D reconstruction pipeline. A crucial aspect in devising a better image matching schema is to adapt the core detector to the recent advancements obtained on the successive steps of the pipeline. On one hand, recent local image descriptors are become robust and able to cope with higher degrees of image deformations and noise (Mishchuk et al.,

\*\*Corresponding author: Tel.: +39-091-23891124  
e-mail: [fabio.bellavia@unipa.it](mailto:fabio.bellavia@unipa.it) (Fabio Bellavia)

2017). On the other hand, the last matching strategies exploiting local spatial constraints (Bellavia, 2022) and those based on the RANdom SAmple Consensus (RANSAC, Fischler and Bolles (1981)) to exploit global model constraints according to scene geometry (Cavalli et al., 2020), have been shown to tolerate much better the presence of outlier matches. Moreover, the increased computational power offered by recent GPU allows to process in parallel more keypoints and matches in a reasonable time depending on the application. It is not a coincidence that the blob-like Difference-of-Gaussian (DoG) keypoints of the Scale Invariant Feature Transform (SIFT) detector (Lowe, 2004), “unchained” to output more keypoints by an appropriate removal of the setup thresholds, have achieved among the best results in the previous Image Matching Challenge<sup>1</sup> (IMC2020, Jin et al. (2020)).

This paper presents HarrisZ<sup>+</sup>, an upgrade to the corner-based HarrisZ detector (Bellavia et al., 2011) for next-gen image matching pipelines. The aim of HarrisZ<sup>+</sup> is to provide more keypoints, better localized and distributed over the image, yet characterized by high repeatability and discriminability. Unlike the case of the unchained SIFT keypoints, HarrisZ<sup>+</sup> does not only consists of a simple tuning of the setup parameters so as to obtain more keypoints as output, but introduces further refinements to the selection criteria delineated by HarrisZ. According to the last Image Matching Challenge<sup>2</sup> (IMC2021) and SimLocMatch Image Matching contest<sup>3</sup>, the image matching pipelines based on HarrisZ<sup>+</sup> obtain state-of-art results among the classic image matching pipelines, closely following the more recent fully deep end-to-end trainable approaches. This proves that there is still a valuable margin of improvement towards the research of classic image matching methods.

The rest of the paper is organized as follows: Sec. 2 presents the related work, Sec. 3 describes the proposed HarrisZ<sup>+</sup> updates after introducing the original HarrisZ detector, and Sec. 4 reports the experimental evaluation. Conclusion and future work are outlined in Sec. 5.

## 2. Related work

Corners and blobs are the two principal kinds of keypoints generally recognized in the literature (Szeliski, 2021), whose the Harris (Harris and Stephens, 1988) and SIFT (Lowe, 2004) detectors may respectively represent the most popular and effective extractors. Corners usually tend to be identified with junctions and blobs with homogeneous image regions, but in practice this distinction only holds for simple synthetic images and not for complex, natural images. The Harris detector is based on the filter response to a function of the eigenvalues of the autocorrelation matrix of the image intensity gradient. Other operators based on the autocorrelation matrix eigenvalues have been proposed to extract corners (Förstner, 1986; Shi

and Tomasi, 1994), but the one defining the Harris corner detector is maybe the most commonly employed. The Hessian matrix has been also used to design keypoint detectors according to its determinant (Beaudet, 1978). Differently from the Harris corner detector, the SIFT detector is based on the filter response to DoG, as approximation of the Laplacian of Gaussian (LoG), which corresponds to the trace of the Hessian matrix. Filter-based detectors are generally invariant to illumination changes since keypoints are selected as local maxima of the filter response. Nevertheless, in practice, an absolute threshold on the filter response is generally set to suppress the detection of spurious noise as keypoints, which can decrease detector capability in case of low-contrast images. Pre-processing the input image by contrast enhancement techniques can be helpful in this situation (Lecca et al., 2020), notwithstanding that noise gets enhanced too. Concerning the robustness to geometric image transformations, in their original definition, detectors are generally only rotational invariant, while the scale depends on the filter windows size. Multi-scale approaches have been devised to solve scale issues, as well as to improve the robustness of the detector to noise (Mikolajczyk and Schmid, 2004). The Gaussian scale-space (Lindeberg, 1994) is generally employed for this aim, yet non-linear scale-space definitions have been successfully applied too, as for the KAZE features (Alcantarilla et al., 2012). The keypoint scale is generally selected according to the maximum filter response among the scales, where the Laplacian filter is a common choice. Additional robustness can be required in case of severe viewpoint changes, which is generally achieved by upgrading the scale and orientation associated with the keypoint to a local affine transformation approximating the unknown original perspective one. The affine transformation can be obtained locally and iteratively (Mikolajczyk and Schmid, 2004), explicitly computing global synthetic warps of the image (Morel and Yu, 2009), but also by implicitly defining the keypoint filter response as a tensor (Zhang and Sun, 2020). This geometric information characterizes the region around the keypoint and is used to generate the local image patch according to a reference system upon which to compute the associated keypoint descriptor. Recent state-of-the-art solutions based on deep learning such as AffNet (Mishkin et al., 2018) only require to provide the keypoint scale and location to get affine-normalized patches.

Fast and efficient keypoint extraction may be required by real-time applications on when processing a huge amount of images. One of the two main approaches in this sense is to discretize the kernel to allow a fast computation of the filter response thanks to the integral images. This solution was originally employed by the Speeded-Up Robust Feature (SURF) detector (Bay et al., 2008) as alternative to SIFT. The other solution is to characterize the keypoint according to a comparison of the keypoint central pixel with respect to the other pixels in the local window. The Features from Accelerated Segment Test (FAST) detector (Rosten et al., 2010) is maybe the most known so-

<sup>1</sup><https://www.cs.ubc.ca/research/image-matching-challenge/2020/>

<sup>2</sup><https://www.cs.ubc.ca/research/image-matching-challenge/>

<sup>3</sup><https://simlocmatch.com/>

lution in this sense. In addition, FAST makes use of decision trees to further speed-up the computation. Worth to be mentioned among the handcrafted keypoint detector is the Maximally Stable Extremal Region (MSER) detector (Matas et al., 2002), which extracts blob-like structures that remain almost stable in a region growing process.

With the partial exception of FAST, machine learning has intervened into the keypoint detector design only recently. After some initial attempts (Verdie et al., 2015), excluding the deep Key.Net detector (Barroso-Laguna et al., 2019) which employs both handcrafted and learned filter layers, deep detectors only appeared as components of end-to-end image matching architectures. The structure of deep detectors generally combine in order convolutional and max pooling layers followed by non-linear activation functions, which strongly resemble the structure of filter-based handcrafted detectors, with respectively a filter definition, the local maxima selection of its response and a thresholding on them. The Learned Invariant Feature Transform (LIFT) network (Yi et al., 2016) is the first example of a complete end-to-end deep image matching architecture. Due to its complexity, LIFT cannot be trained as a whole from scratch and the training process requires handcrafted keypoints as bootstrap. An alternative solution was proposed in SuperPoint (DeTone et al., 2018), where the training process becomes self-supervised thanks to a pre-training phase with synthetic images and the homographic adaptation. Unlike LIFT which is patch-based, SuperPoint relies on a fully convolutional network to process the whole image as input. The Detect-and-Describe (D2-Net, Dusmanu et al. (2019)) and the Describe-to-Detect (D2D, Tian et al. (2020)) networks are other end-to-end local feature architectures, which operate in a different way to the standard computational flow where it is the descriptor that must be adapted to the keypoint detector. Specifically, D2-Net defines dense feature maps that simultaneously serve to derive the keypoints and their descriptors, while in D2D keypoints must be adapted to the information content of the descriptors. The DIScrete Keypoints (DISK) network (Tyszkiewicz et al., 2020) leverages the principles of reinforcement learning to ease the complexity of the end-to-end training caused by the sparseness of the keypoints and achieved state-of-the-art results. More recently, the integration of the keypoint spatial constraints into the network design has greatly improved the final matching output. This has been done implementing attentional graph neural networks in SuperGlue (Sarlin et al., 2019), using coarse-to-fine schemas with the Local Feature TRAsformer (LoFTR, Sun et al. (2021)), or considering dense descriptor approaches not employing any sort of keypoints (Truong et al., 2021).

### 3. HarrisZ<sup>+</sup>

#### 3.1. From HarrisZ

The original HarrisZ (Bellavia et al., 2011) extracts corners from an image, represented as a matrix  $I \in \mathbb{R}^{m \times n}$ , at different scales of the Gaussian scale-space. Defining

the horizontal and vertical gradient derivatives of  $I$  as the image convolution with the central difference kernel  $K = [-1 \ 0 \ 1]$ , i.e.

$$I_x = I * K, \quad \text{and} \quad I_y = I * K^\top \quad (1)$$

where  $*$  indicates the convolution, the scale-space derivatives are defined as (Mikołajczyk and Schmid, 2004)

$$I_{\sigma,x} = G_\sigma * I_x \quad \text{and} \quad I_{\sigma,y} = G_\sigma * I_y \quad (2)$$

where  $\sigma$  denotes the differentiation scale which identifies the current working scale and  $G_\sigma$  is a Gaussian kernel with mean zero and standard deviation  $\sigma$ . Unlike the standard Harris corner detector (Harris and Stephens, 1988), before computing the autocorrelation matrix both the scale-space derivatives get enhanced by a pixel-wise multiplication according to a smoothed raw edge mask  $M_\sigma$

$$E_{\sigma,x} = I_{\sigma,x} M_\sigma \quad \text{and} \quad E_{\sigma,y} = I_{\sigma,y} M_\sigma \quad (3)$$

where  $M_\sigma$  is computed from the gradient magnitude

$$I_{\sigma,m} = \left( I_{\sigma,x}^2 + I_{\sigma,y}^2 \right)^{\frac{1}{2}} \quad (4)$$

by globally threshold on the mean  $\bar{I}_{\sigma,m}$  computed on the whole  $I_{\sigma,m}$ , i.e.

$$M_\sigma = G_\sigma * \mathcal{V}(I_{\sigma,m} > \bar{I}_{\sigma,m}) \quad (5)$$

where  $\mathcal{V}(\cdot)$  is the indicator function. The enhanced derivatives decrease the noise while strengthening the edges according to the working scale  $\sigma$ , providing a sort of non-linear scale-space. The autocorrelation matrix at the pixel  $\mathbf{x}$  is then computed as

$$\mu_\sigma(\mathbf{x}) = \begin{bmatrix} G_{s\sigma} * E_{\sigma,x}^2(\mathbf{x}) & G_{s\sigma} * E_{\sigma,x}(\mathbf{x}) E_{\sigma,y}(\mathbf{x}) \\ G_{s\sigma} * E_{\sigma,x}(\mathbf{x}) E_{\sigma,y}(\mathbf{x}) & G_{s\sigma} * E_{\sigma,y}^2(\mathbf{x}) \end{bmatrix} \quad (6)$$

where  $s\sigma$  defines the integration scale, i.e. the window size of the filter, on the basis of the differentiation scale  $\sigma$  and the constant  $s = 2^{\frac{1}{2}}$ . Actually, in the computation of the corner filter response  $H_\sigma$ , only the determinant  $D_\sigma$  and the trace  $T_\sigma$  of  $\mu_\sigma$  are involved, respectively equal to the product and the sum of the eigenvalues of  $\mu_\sigma$ . These can be directly computed by simple pixel-wise operations on Gaussian-smoothed maps as

$$D_\sigma = (G_{s\sigma} * E_{\sigma,x}^2)(G_{s\sigma} * E_{\sigma,y}^2) - (G_{s\sigma} * (E_{\sigma,x} E_{\sigma,y}))^2 \quad (7)$$

$$T_\sigma = (G_{s\sigma} * E_{\sigma,x}^2) + (G_{s\sigma} * E_{\sigma,y}^2) \quad (8)$$

In the original definition given by Harris and Stephens, the corner filter response is given by

$$\hat{H}_\sigma = D_\sigma - c T_\sigma^2 \quad (9)$$

where  $c$  is a user-given coefficient. The HarrisZ filter response  $H_\sigma$  replaces  $c$  with an implicit adaptive reformulation according to the z-score normalization  $\mathcal{Z}(\cdot)$  acting globally on the whole image

$$H_\sigma = \mathcal{Z}(D_\sigma) - \mathcal{Z}(T_\sigma^2) \quad (10)$$

where

$$\mathcal{Z}(Q) = \frac{Q - \bar{Q}}{\varsigma_Q} \quad (11)$$

being  $Q$  a generic image with mean  $\bar{Q}$  and standard deviation  $\varsigma_Q$ . According to the statistical properties of  $H_\sigma$  considering the whole input image  $I$ , a pixel  $\mathbf{x}$  can be roughly classified as a flat region if  $H_\sigma(\mathbf{x})$  is almost 0, or otherwise as an edge or as a corner if  $H_\sigma(\mathbf{x})$  is respectively lower or higher than 0. Moreover, a corner must lie onto the edge mask  $M_\sigma$ . These two conditions can be coded as the binary mask  $C_\sigma$

$$C_\sigma = \mathcal{V}(H_\sigma > 0) \mathcal{V}(M_\sigma > 0.31) \quad (12)$$

where multiplication is intended pixel-wise and the 0.31 value is introduced to take into account the current scale in the binarization of  $M_\sigma$ , whose formal derivation can be found in the original manuscript. The binary mask  $C_\sigma$  provides the initial set of corner candidates according to the input image instead of requiring a user-defined threshold which generally needs to be adjusted according to the input, as happens for the original Harris detector.  $C_\sigma$  provides a sort of implicit adaptive contrast enhancement of the input image since both  $H_\sigma$  and  $M_\sigma$  rely on global statistics for their definitions. Note that recent state-of-the-art deep architectures use masks to filter keypoints with the aim of removing disturbing elements of the scene such as sky or people<sup>2</sup>.

Finally, corners corresponding to  $C_\sigma$  are selected so that only local maxima on  $H_\sigma$  will survive. Differently from the standard multi-scale approach employed for instance in the Harris-affine detector, where the convolution kernels are maintained fixed while the image is down-scaled, HarrisZ keeps the input image size fixed while increasing the kernel size, as for SURF (Bay et al., 2008), in order to increase the detector accuracy. The final set of keypoints for the scale  $\sigma$  is selected by sorting first the  $C_\sigma$  pixel locations according to the decreasing values of  $H_\sigma$  and then, in order, greedily removing a candidate keypoint if its distance from the already selected keypoints is less than  $\lceil 3\sigma \rceil$ . This strategy can resemble in some aspect the Adaptive Non-Maximal Suppression (ANMS, Brown et al. (2005)).

Keypoint localization precision is then increased with parabolic sub-pixel precision on  $H_\sigma$ . The entire process is repeated for different scales

$$\sigma \in \{2^{\frac{i}{2}} : 3 \leq i \leq 8\} \quad (13)$$

Corners are filtered one last time according to the ratio of the corresponding autocorrelation matrix eigenvalues, which must be higher than 0.75 to further exclude accidental detected edges. Finally, corners are considered altogether among the different scales and sorted first by the decreasing scale values index  $i$  and then by  $H_\sigma$  if the maximum number of allowed output keypoints is exceeded.

### 3.2. To HarrisZ<sup>+</sup>

HarrisZ<sup>+</sup> differs from the original HarrisZ by the fine tuning of the parameters in combination with fine imple-

mentation changes: the devil is in the details. These updates are strongly related to the advances introduced in the successive steps of the matching pipeline, each one leading to state-of-the-art improvements. In particular:

- i AffNet (Mishkin et al., 2018) provides a better patch normalization than the canonical one of SIFT, which gives only the scale<sup>4</sup>, and it also performs better with respect to the affine transformation estimated by the autocorrelation matrix (Mikolajczyk and Schmid, 2004).
- ii The last version of HardNet (Pultar, 2020) provides a more discriminant and robust keypoint descriptor than SIFT.
- iii Blob matching strategy with the inclusion of spatial constraints based on the Delaunay Triangulation Matching (DTM, Bellavia (2022)) retains more correct matches while discarding wrong matches better than the simple Nearest Neighbour Ratio (NNR) strategy (Lowe, 2004).
- iv Filtering matches by multiple local RANSAC has proven to be useful in removing wrong matches with the Adaptive Locally-Affine Matching (AdaLAM, Cavalli et al. (2020)).
- v The Degenerate SAmple Consensus (Degenerate SAC, (Chum et al., 2005)) checks for degenerate model configurations in RANSAC and obtains better results than the standalone RANSAC.

Furthermore, the computational power and memory availability of recent hardware greatly increased in the last decade. On this basis, HarrisZ<sup>+</sup> adds incrementally the following updates and upgrades with respect to HarrisZ:

1. In the original HarrisZ, keypoints are scored first by their decreasing scale indexes  $i$  and then by  $H_\sigma$  so as to extract the most robust keypoint since higher scales imply wider patch support regions. Nevertheless, the keypoint localization accuracy at high scales degrades. The potential loss in the discriminative power of the local keypoint can be compensated by the choice of a better keypoint descriptor, so HarrisZ<sup>+</sup> scores keypoints first by  $H_\sigma$  and then by the decreasing scale index  $i$ . Furthermore, the First Geometric Inconsistent Nearest Neighbor (FGINN, Mishkin et al. (2015)) is employed instead of NNR for matching, since FGINN can handle better Jin et al. (2020) than NNR matches in case of spatially close keypoints, which are promoted by the new keypoint ranking. Specifically, if a keypoint  $k_1$  on an image  $I_1$  matches with a keypoint  $k_2$  on the other image  $I_2$ , this means that their descriptor distance is the minimum  $d$  among the distances of remaining keypoints in  $I_2$  with  $k_1$ . In order to filter out unreliable matches, NNR normalizes  $d$  by the minimum distance  $d'$  between  $k_1$  and the remaining keypoints in  $I_2$  excluding  $k_2$  and thresholds this

<sup>4</sup>patch orientation is not taken into account since only almost aligned images are considered, see Sec.4

ratio at a value typically in  $[0.8, 0.95]$ . This greatly increases precision among the tentative matches, but often reduces recall by filtering out correct matches. FGINN restricts the minimum  $d'$  only to keypoints far more than 10 px with respect to  $k_2$ , so closely detected points do not hurt each others.

- For ranking keypoints in the case where a constraint on their maximum number has been imposed, a strategy inspired by the greedy local maxima selection of Sec. 3.1 is employed to better distribute the corners over the images, avoiding image regions with clusters of keypoints which can reasonably be associated to the same true corners. If the input image is of size  $m \times n$  px and a constraint of having no more than  $k$  keypoints is imposed, the minimum required distance  $q$  between two keypoints is set as the diameter of a circle with area  $\frac{mn}{k/2}$ , i.e.

$$q = (2^3 mn)^{\frac{1}{2}} / \pi k \quad (14)$$

Keypoints are then sorted as indicated above, and the greedy local maxima selection using  $q$  as distance threshold is performed. If less than  $k$  keypoints survive, the process is repeated again on the discarded keypoints. Notice that only roughly  $\frac{k}{2}$  keypoints are expected to be selected in the first iteration, so it is more likely that in the second iteration some keypoints close to those found in the first iteration will be chosen. According to preliminary evaluations, this strategy leads to better matching results, probably due to the fact of allowing to some extent closer keypoint, which generally have some variation within their descriptors, introduces a further chance to match the right descriptors.

- The scale index  $i$  for  $\sigma$  in Eq. 13 is set to  $0 \leq i \leq 4$ . The original scale index range  $3 \leq i \leq 8$  was chosen to limit to about 2000 keypoints or less the detector output for common input image resolution ( $1024 \times 768$  px), but also to extract the most robust keypoints, since higher scales imply wider patch support regions, but at the expense of the keypoint localization accuracy. The new range can extract more than 8000 keypoints, which according to recent evaluations greatly increase the matching power (Jin et al., 2020). Moreover, at finer scale the keypoint localization accuracy is improved. Again, the potential loss in the of discriminative power of the patches is compensated by the choice a better keypoint descriptor.
- For color images, the original HarrisZ detector uses the common approach of performing a grayscale conversion of the input using only the luminance channel (Gonzalez and Woods, 2008). Another popular choice for graylevel conversion is the use of the value channel of the HSV decomposition, corresponding to take the maximum value over the RGB channels. This further grayscale conversion can highlight different image structures, including edges. According to this observation, the edge mask  $M_\sigma$  in Eq. 5 is computed

from further enhanced gradient derivatives  $I_x^\star$  and  $I_y^\star$  of  $I$  which take into account both the luminance  $L$  and value  $V$  channels in the grayscale conversion of the input image  $I$ . Specifically, the base horizontal derivatives of  $L$  and  $V$  are computed as for Eq. 1

$$L_x = L * K \quad \text{and} \quad V_x = V * K \quad (15)$$

so that the maximum absolute gradient value at each pixel location  $\mathbf{x}$  is chosen

$$I_x^\star(\mathbf{x}) = \begin{cases} L_x(\mathbf{x}) & \text{if } |L_x(\mathbf{x})| > |V_x(\mathbf{x})| \\ V_x(\mathbf{x}) & \text{otherwise} \end{cases} \quad (16)$$

and likewise for  $I_y^\star$  from the vertical derivatives  $L_y$  and  $V_y$ . As shown in Fig. 1,  $I_x^\star$  and  $I_y^\star$  produce a better edge mask  $M_\sigma$  but, according to our preliminary experiments, they should not be used for computing the filter response  $H_\sigma$  since they lead to worse results.

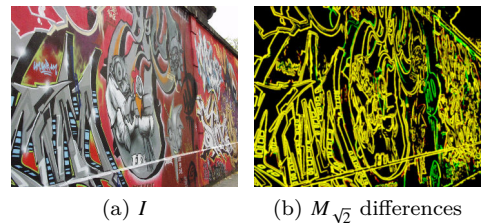


Fig. 1: Computation of the mask  $M_\sigma$  for a color image  $I$ . Mask portions retained in any case are in yellow, while mask regions considered only with or without the integration of the HSV value channel are respectively in green and red (see text for details, best viewed in color and zoomed in).

- In the actual setup, keypoint patches at scale index  $i = 0$  will get a support region of  $25 \times 25$  px, while HardNet descriptor input patches are  $32 \times 32$  px (Mishchuk et al., 2017). Upsampling would be required in this case, which can introduce noise and spurious details affecting the descriptor robustness. In order to avoid patch upsampling, the final scale associated to the keypoints extracted at  $i = 0$  at the end of the computation is replaced with the scale value associated to  $i = 1$ , which provides  $35 \times 35$  px patches. Patch downsampling does not negatively affect the HardNet descriptor.
- As keypoints for  $i = 0, 1$  get the same final scale  $\sigma$ , i.e. both scale indexes output  $35 \times 35$  px patches, local maxima on their union is considering following again the greedy approach with a distance of 1 px. This avoids almost-duplicated keypoint patches.
- Before being processed as described in Sec. 3.1, the input image for scales with indexes  $i = 0, 1$  is doubled using Lanczos interpolation so as the scale value  $\sigma$  to reflect the image size adjustment. Extracted keypoints are then reported back to the original reference system by halving the coordinate and scale values. This modification aims at introducing sub-pixel quality details and hence to finer characterizing keypoints, also under the observation that discretized Gaussian



kernels with standard deviations around 1 pixel for practical purposes provide basically the same convolution output ( $\sigma = 1, 1.4$  px for  $i = 0, 1$ ). Besides Lanczos, the bilinear and bicubic interpolations were considered, but they provided worse results.

Figure 2 shows the differences between HarrisZ and HarrisZ<sup>+</sup>, underling the relations between scales and localization accuracy, as well as the different distribution of the corners on the image.

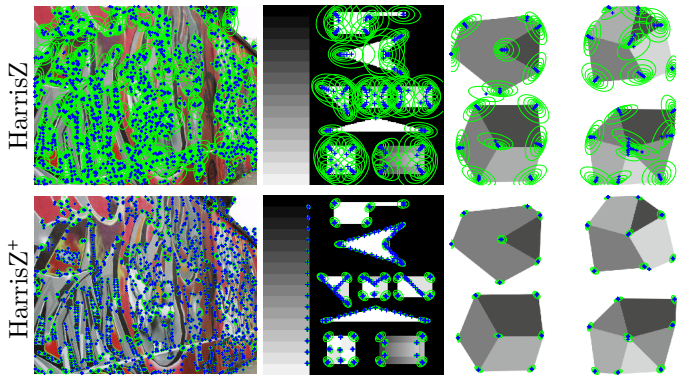


Fig. 2: Comparison between HarrisZ and HarrisZ<sup>+</sup> on example images, for each keypoint the ellipses associated to the affine transformation determined by  $\mu_\sigma$  is shown in green (see text for details, best viewed in color and zoomed in).

## 4. Evaluation

### 4.1. Setup

The evaluation relies on the IMC and SimLocMatch contests.

The IMC benchmark (Jin et al., 2020) evaluates HarrisZ<sup>+</sup> as the upstream module of the image matching pipeline steps i-v described in Sec. 3.2, considering the camera pose error obtained on different scenes. In detail, the error is defined in terms of the mean Average Accuracy (mAA), which takes into account the error achieved at the end of the matching pipeline after the stereo or multiview camera pose estimation. mAA is computed by integrating the maximum angular error between the rotation and translation vectors of the final estimated fundamental matrix up to a threshold of  $10^\circ$  for each tested image pairs of the scene. For this aim, reliable pseudo ground-truths obtained by supersets of the image sequences are employed. mAA correlates well with the matching score, but not with repeatability, as the latter is sensitive to the number of keypoints and does not generally agree with the desired and expected results. The IMC benchmark datasets are Phototourism, PragueParks and GoogleUrban, all made up of real images and including validation and test sets. Overall, the complexity from one dataset to another is increasing: each Phototourism scene consists of a main large-scale foreground object almost taken in frontal position, PragueParks includes small-scale scenes with multiple complex objects and higher camera pose variations, and GoogleUrban presents challenge middle-scale urban scene environments with a low level of image overlap. The

test sets of the IMC datasets contain respectively 10, 3 and 17 image sequences. Except GoogleUrban, where each image sequence contains 64 or 75 images, other datasets include 100 images for each sequence. Besides the three Phototourism validation sequences employed to check and adjust HarrisZ<sup>+</sup> steps, the other validation sets were only used to setup the best DegenSAC threshold and not to fine tune any network parameter in the pipeline, unlike other competitors. Lastly, following the IMC protocol, evaluation is done by restricting the maximum number of allowed keypoints to 2048 (2K) and 8000 (8K).

Differently from IMC, the SimLocMatch benchmark relies on rendered images from synthetic scenes. In this way not only the poses are available for validating the results, but also the exact depth maps so that the evaluation can be directly focused on matches instead of using indirectly the camera pose estimation. The SimLocMatch dataset includes seven scenes representing man-made outdoor and indoor environments for a total of about 10000 images, with different simulated lighting conditions. In the evaluation, matches from roughly 80000 image pairs are checked and, unlike IMC, these include negative image pairs, i.e. with no matches between the images. As evaluation criterion, the SimLocMatch benchmark considers the matching score and the number of wrong matches per image pair, respectively having higher and lower values in the case of better methods.

Notice that both benchmarks only consider almost upright images, i.e. with no relevant relative rotation, which is the most common user application scenario and represents the standard training data for deep architecture. As rotation invariance decrease matchability, the former is disabled for a fair comparison in all methods.

The Matlab code for both HarrisZ and HarrisZ<sup>+</sup> corner detectors is freely available to download, as well as the full image matching pipeline implementation<sup>5</sup> which makes use of the code available from the respective authors or through the Kornia library (Riba et al., 2020).

### 4.2. Results

Table 1 shows the results on the IMC2021 Phototourism validation set obtained by incrementally incorporating the additions presented in Sec. 3.2 to get HarrisZ<sup>+</sup> from HarrisZ, the final HarrisZ<sup>+</sup> corresponds to the last row. For this evaluation, the successive steps of the pipeline consider the AffNet for patch extraction, HardNet8 for computing the descriptors and DegenSAC to estimate the pose (listed as i, ii and v in Sec. 3.2). Spatial local filtering (iii and iv) is excluded to better evaluate the keypoint extraction process. This “ablation study” only considers the stereo setup, excluding multiview pose optimization by bundle adjustment with COLMAP (Schönberger and Frahm, 2016). Each addition from HarrisZ can be applied independently with moderate improvements, however only when taken altogether the proposed changes get an overall significant gain. Specifically, according to the results,

<sup>5</sup><https://sites.google.com/view/fbellavia/>

better localized keypoints provided by lower scales (listed as 1 and 3 in Sec. 3.2) are more critical in 8K setup. On the contrary, when the keypoint budget is reduced to 2K, providing an uniform distribution of keypoints among the image (2) becomes relevant. The remaining extensions (4-7) serve mainly as refinements and although each of them only introduces relatively small improvements, the final mAA increase globally of roughly 0.1 and 0.05 respectively for the 2K and 8K setups.

Table 1: Incremental additions towards HarrisZ<sup>+</sup> and corresponding mAA@10° on IMC2021 Phototourism validation set. The # column refers to the last change described in Sec. 3.1 sequentially included from the original HarrisZ, the final HarrisZ<sup>+</sup> corresponds to the last row (see text for details)

Incremental additions	#	mAA@10°	
		2K	8K
Original HarrisZ	-	0.4619	0.5980
$H_\sigma$ ranking + FGINN	1	0.4955	0.6209
Uniform distribution of corners	2	0.5335	0.6203
Scale indexes $i = 0, \dots, 4$	3	0.5410	0.6422
Color image mask $M_\sigma$	4	0.5390	0.6515
Output scale $\sigma$ adjustment for $i = 0$	5	0.5493	0.6520
Duplicates removal for $i = 0, 1$	6	0.5636	0.6522
Doubling input image size for $i = 0, 1$	7	0.5681	0.6565

Table 2 reports the full mAA results for both the stereo and multiview setups of the last IMC2021, which adds two further datasets with respect to the Phototourism one presented in the previous IMC2020. The HarrisZ<sup>+</sup> pipeline including also blob matching plus DTM spatial filtering (iii in Sec. 3.2) with or without AdaLAM (iv) is evaluated. All the pipelines include DegenSAC (v) as final step, followed by bundle adjustment in case of the multiview setup. The standard upright SIFT matching is included for reference, while the pipeline that replaces HarrisZ<sup>+</sup> with the unchained SIFT, i.e. the DoG detector, obtained results among the state-of-the-art in the previous IMC2020 together with DISK. The remaining entries obtained the best results in IMC2021 and employ in addition learned matching strategies. Among the non-learned matching approaches, the pipelines including HarrisZ<sup>+</sup> provide the best results on each dataset with the exception of Phototourism, where results are almost aligned, with some advantages of DISK on the 2K keypoint restriction. Among the compared handcrafted keypoint detectors, HarrisZ<sup>+</sup> obtains the best results followed by DoG and then SIFT, implying the goodness of the proposed approach, while adding spatial filtering to the pipeline seems to only improve the results in the multiview setup. Excluding the GoogleUrban dataset and the Phototourism stereo setup, when comparing the pipelines based on HarrisZ<sup>+</sup> with respect to the best ones relying on learned matching strategies, results are quite comparable. The performance differences are mostly noticeable within the most challenging GoogleUrban dataset, and are probably due to the ability of these architectures to employ, more or less implicitly, a coarse-to-fine image matching strategy, missing in other approaches.

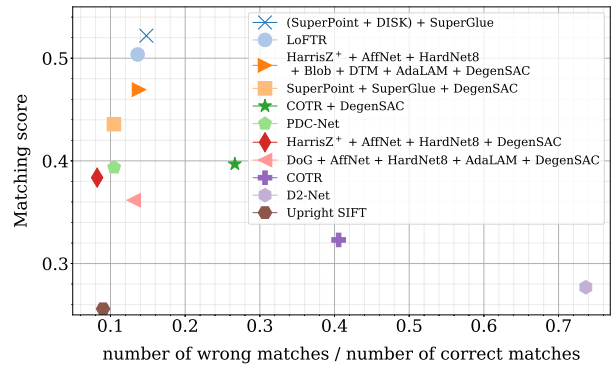


Fig. 3: SimLocMatch evaluation results (see text for details, best viewed in color and zoomed in).

Figure 3 reports the SimLocMatch results. HarrisZ<sup>+</sup> and DoG are limited to 8K keypoints, while the CORrespondence TRansformer (COTR, Jiang et al. (2021)) and the Probabilistic Dense Correspondence Network (PDC-Net, Truong et al. (2021)) are included for completeness. Confirming the IMC2021 results, HarrisZ<sup>+</sup> pipeline completed with spatial local match filtering is the only pipeline able to obtain results very close to those of the approaches employing matching networks based on transformers, i.e. LoFTR and (SuperPoint and Disk) + SuperGlue. In particular, it can be noted that besides achieving a high matching score, also the number of wrongly output matches, normalized by the number of inliers for visualization purposes, is low for HarrisZ<sup>+</sup>.

Concerning running times, non-optimized HarrisZ<sup>+</sup> Matlab code requires on average one second to run on a Intel i9-10900K system with 64 GB RAM and no GPU in case of images with resolutions going from  $640 \times 480$  to  $1900 \times 1200$  px, which is roughly five times the original HarrisZ runtime. The bottleneck is the doubled image size for the first two scale indexes. Moreover, HarrisZ<sup>+</sup> keypoint uniform ranking is a global process performed at the end of the computation, which can strongly accentuate the computational times in case of a high number of candidate keypoints, as for high-textured images. Nevertheless, these running times are still reasonable for off-line applications.

As final consideration, classical modular pipeline design, as that HarrisZ<sup>+</sup> was designed for, can still offer practical advantages with respect to modern end-to-end deep design for specific applications in terms computational cost or tuning. Adding rotation invariance to the HarrisZ<sup>+</sup> pipeline only requires to simply plug OriNet (Mishkin et al., 2018) or similar modules before AffNet, while both SuperGlue and LoFTR would require retraining and maybe to add some layers. In any case retrain of a whole end-to-end matching network needs a more expensive system configuration than that mentioned above (with the inclusion of a consumer-grade GPU). On this system, HarrisZ<sup>+</sup> pipeline was also tested to work with high resolution images up to  $9000 \times 6732$  px, such those of aerial photogrammetric surveys, while LoFTR was unable

Table 2: mAA@10° results on IMC2021 (see text for details).

Method	Phototourism		PragueParks		GoogleUrban		Total		
	2K	8K	2K	8K	2K	8K	2K	8K	
Stereo	HarrisZ <sup>+</sup> + AffNet + HardNet8	0.4753	0.5578	0.6358	0.7368	0.3252	0.3567	0.4788	0.5504
	+ Blob + DTM	0.4449	0.5515	0.6176	0.7387	0.2981	0.3343	0.4535	0.5415
	+ Blob + DTM + AdaLAM	0.4369	0.5521	0.5918	0.7423	0.2963	0.3327	0.4417	0.5424
	Upright SIFT	0.3827	0.5122	0.4136	0.5369	0.2542	0.2693	0.3501	0.4394
	DoG + AffNet + HardNet8	n/a	0.5573	n/a	0.5977	n/a	0.3009	n/a	0.4853
	DISK	0.5121	0.5583	0.4589	0.5622	0.2763	0.3295	0.4157	0.4833
	LoFTR	n/a	0.6090	n/a	0.7546	n/a	0.4060	n/a	0.5898
	(SuperPoint and DISK) + SuperGlue	n/a	0.6397	n/a	0.8070	n/a	0.4395	n/a	0.6285
	HarrisZ <sup>+</sup> + AffNet + HardNet8	0.6786	0.7367	0.4676	0.4729	0.1602	0.2025	0.4354	0.4707
	+ Blob + DTM	0.7046	0.7606	0.4583	0.4843	0.1625	0.2060	0.4418	0.4836
+ Blob + DTM + AdaLAM	0.7132	0.7580	0.4618	0.4711	0.1521	0.2084	0.4423	0.4791	
Multiview	Upright SIFT	0.5545	0.6850	0.3608	0.4810	0.0520	0.0902	0.3224	0.4187
	DoG + AffNet + HardNet8	n/a	0.7269	n/a	0.4670	n/a	0.1643	n/a	0.4527
	DISK	0.7296	0.7445	0.4357	0.4590	0.1275	0.1833	0.4309	0.4623
	LoFTR	n/a	0.7610	n/a	0.4712	n/a	0.3023	n/a	0.5115
	(SuperPoint and DISK) + SuperGlue	n/a	0.7857	n/a	0.4988	n/a	0.3374	n/a	0.5406

to process  $2000 \times 1500$  px input images on a consumer-grade PC either on GPU or CPU.

## 5. Conclusions and future work

This paper presents HarrisZ<sup>+</sup> corner detector as an upgrade of HarrisZ to be employed with next-gen image matching pipeline. By introducing several changes, which by themselves only slightly affect the final output, HarrisZ<sup>+</sup> is able to achieve state-of-the-art results among handcrafted keypoint detectors, providing the basis of competitive image pipelines with respect to the recent end-to-end deep architectures exploiting coarse-to-fine matching strategies thorough transformers. The key idea that guided the HarrisZ<sup>+</sup> design is the need of a synergic optimization between the different modules of the pipeline to globally improve the final results, so as to compensate weakness and boost strength points of the whole pipeline.

It is also authors’ opinion that research on classical matching methods is still appropriate and effective despite of the overwhelming presence of deep methods. As a further contribution of this paper, on one hand the results achieved through HarrisZ<sup>+</sup> shows that there is still a worthwhile margin of improvement in classic research direction, and on the other hand the gained insight during the development can be employed in turn to design better deep networks as well as to provide more valid training data.

Future work will includes the introduction of handcrafted a coarse-to-fine matching at the end of the pipeline with further optimizations of HarrisZ<sup>+</sup> to better adapt to these changes and to improve the computational efficiency, investigations about the possibility to apply HarrisZ<sup>+</sup>-like selection to DoG keypoints and about using HarrisZ<sup>+</sup> keypoints with deep matching architectures based on transformers, such as SuperGlue.

## Acknowledgments

F. Bellavia is funded by the Italian Ministry of Education and Research (MIUR) under the program PON

Ricerca e Innovazione 2014-2020, cofunded by the European Social Fund (ESF), CUP B74I18000220006, id. proposta AIM 1875400, linea di attività 2, Area Cultural Heritage. D. Mishkin is supported by OP VVV funded project CZ.02.1.01/0.0/0.0/16\_019/0000765 “Research Center for Informatics” and CTU student grant SGS20/171/OHK3/3T/13.

## References

- Alcantarilla, P.F., Bartoli, A., Davison, A.J., 2012. KAZE features, in: Proceedings of the European Conference on Computer Vision (ECCV).
- Barroso-Laguna, A., Riba, E., Ponsa, D., Mikolajczyk, K., 2019. Key.Net: Keypoint detection by handcrafted and learned CNN filters, in: Proceedings of the International Conference on Computer Vision (ICCV).
- Bay, H., Ess, A., Tuytelaars, T., Gool, L.V., 2008. Surf: Speeded up robust features. *Computer Vision and Image Understanding* 110, 346–359.
- Beaudet, P., 1978. Rotationally invariant image operators, in: International Joint Conference on Pattern Recognition, pp. 578–583.
- Bellavia, F., 2022. SIFT matching by context exposed. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Bellavia, F., Tegolo, D., Valenti, C., 2011. Improving Harris corner selection strategy. *IET Computer Vision*.
- Brown, M., Szeliski, R., Winder, S., 2005. Multi-image matching using multi-scale oriented patches, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 510–517.
- Cavalli, L., Larsson, V., Oswald, M.R., Sattler, T., Pollefeys, M., 2020. AdaLAM: Revisiting handcrafted outlier detection, in: ECCV.
- Chum, O., Werner, T., Matas, J., 2005. Two-View Geometry Estimation Unaffected by a Dominant Plane, in: CVPR.
- DeTone, D., Malisiewicz, T., Rabinovich, A., 2018. Superpoint: Self-supervised interest point detection and description, in: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops.
- Dusmanu, M., Rocco, I., Pajdla, T., Pollefeys, M., Sivic, J., Torii, A., Sattler, T., 2019. D2-Net: A Trainable CNN for Joint Detection and Description of Local Features, in: Proceedings of the 2019 IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
- Fischler, M., Bolles, R., 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis



- and automated cartography. *Communications of the ACM* 24, 381–395.
- Förstner, W., 1986. A feature based correspondence algorithm for image matching. *International Archives of Photogrammetry and Remote Sensing* 26, 150–166.
- Gonzalez, R.C., Woods, R.E., 2008. *Digital image processing*. Prentice Hall.
- Harris, C., Stephens, M., 1988. A combined corner and edge detector, in: *Proceedings of the 4th Alvey Vision Conference*, pp. 147–151.
- Jiang, W., Trulls, E., Hosang, J., Tagliasacchi, A., Yi, K.M., 2021. COTR: Correspondence transformer for matching across images, in: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Jin, Y., Mishkin, D., Mishchuk, A., Matas, J., Fua, P., Yi, K.M., Trulls, E., 2020. Image matching across wide baselines: From paper to practice, in: *arXiv*.
- Lecca, M., Torresani, A., F. Remondino, 2020. Comprehensive evaluation of image enhancement for unsupervised image description and matching. *IET Image Processing*.
- Lindeberg, T., 1994. *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers.
- Lowe, D., 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60, 91–110.
- Matas, J., Chum, O., Urban, M., Pajdla, T., 2002. Robust wide baseline stereo from maximally stable extremal regions, in: *British Machine Vision Conference*, pp. 384–393.
- Mikolajczyk, K., Schmid, C., 2004. Scale & affine invariant interest point detectors. *International Journal of Computer Vision* 60, 63–86.
- Mishchuk, A., Mishkin, D., Radenovic, F., Matas, J., 2017. Working Hard to Know Your Neighbor’s Margins: Local Descriptor Learning Loss, in: *NeurIPS*.
- Mishkin, D., Matas, J., Perdoch, M., 2015. MODS: Fast and robust method for two-view matching. *Computer Vision and Image Understanding* 141, 81 – 93.
- Mishkin, D., Radenovic, F., Matas, J., 2018. Repeatability is Not Enough: Learning Affine Regions via Discriminability, in: *ECCV*.
- Morel, J., Yu, G., 2009. ASIFT: A new framework for fully affine invariant image comparison. *SIAM Journal on Imaging Sciences* 2, 438–469.
- Mur-Artal, R., Montiel, J., Tardos, J., 2015. ORB-SLAM: a versatile and accurate monocular slam system. *IEEE Transactions on Robotics* 31, 1147–1163.
- Pultar, M., 2020. Improving the HardNet descriptor. *arXiv ePrint* 2007.09699.
- Riba, E., Mishkin, D., Ponsa, D., Rublee, E., Bradski, G., 2020. Kornia: an open source differentiable computer vision library for pytorch, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.
- Rosten, E., Porter, R., Drummond, T., 2010. Faster and better: A machine learning approach to corner detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 105–119.
- Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A., 2019. SuperGlue: Learning feature matching with graph neural networks, in: *arXiv*.
- Schönberger, J.L., Frahm, J.M., 2016. Structure-from-Motion revisited, in: *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Shi, J., Tomasi, C., 1994. Good features to track, in: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 593–600.
- Sun, J., Shen, Z., Wang, Y., Bao, H., Zhou, X., 2021. LoFTR: Detector-free local feature matching with transformers.
- Szeliski, R., 2021. *Computer Vision: Algorithms and Applications*, 2nd edition. Springer.
- Tian, Y., Balntas, V., Ng, T., Laguna, A.B., Demiris, Y., Mikolajczyk, K., 2020. D2D: Keypoint extraction with describe to detect approach, in: *Proceedings of the 15th Asian Conference on Computer Vision (ACCV)*.
- Truong, P., Danelljan, M., Gool, L.V., Timofte, R., 2021. PDC-Net: Learning accurate dense correspondences and when to trust them, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tuytelaars, T., Mikolajczyk, K., 2008. Local invariant feature detectors: a survey. *Foundations and Trends in Computer Graphics and Vision* 3, 177–280.
- Tyszkiewicz, M.J., Fua, P., Trulls, E., 2020. DISK: Learning local features with policy gradient, in: *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS)*.
- Verdie, Y., Yi, K.M., Fua, P., Lepetit, V., 2015. TILDE: A temporally invariant learned detector., in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5279–5288.
- Yi, K.M., Trulls, E., Lepetit, V., Fua, P., 2016. Lift: Learned invariant feature transform, in: *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Zhang, W., Sun, C., 2020. Corner detection using multi-directional structure tensor with multiple scales. *International Journal of Computer Vision*.