Low-Complexity Pruning for Accelerating Corner Detection

Meiqing Wu, Nirmala Ramakrishnan, Siew-Kei Lam and Thambipillai Srikanthan
Centre for High Performance Embedded Systems
Nanyang Technological University
Singapore
{wume0007, rama0056}@e.ntu.edu.sg, {assklam, astsrikan}@ntu.edu.sg

***Abstract*—In this paper, we present a novel and computationally efficient pruning technique to speed up the Shi-Tomasi and Harris corner detectors. The proposed technique quickly prunes non-corners and selects a small corner candidate set by approximating the complex corner measure of Shi-Tomasi and Harris. The actual corner measure is then applied only to the reduced candidate set. Experimental results on the NiOS-II platform show that the proposed technique achieves an average execution time savings of 90% for Shi-Tomasi and 70% for Harris detectors for 500 corners with no loss in accuracy.**

I. INTRODUCTION

Feature detection and extraction is the fundamental step in most image processing and computer vision applications such as image registration, object detection and video tracking. Corners are highly preferred features as they are present in all types of images. Shi-Tomasi [1] and Harris [2] are classic and widely used corner detection algorithms. Recent evaluations in real time applications such as video tracking [3], visual SLAM [4] and robotic navigation [5] have demonstrated that both of these algorithms are also preferred as feature detectors in combination with more complex feature descriptors. However, both the algorithms require a complex corner measure computation for every pixel in the image. This step is highly compute intensive and becomes a bottleneck for real time vision tasks such as high frame rate video processing [6]. Reducing the computational complexity of corner detection algorithms is essential in low cost embedded systems that do not support micro-processors with sophisticated arithmetic units or hardware accelerators.

In this work, we propose a novel low complexity pruning technique that removes the non-corners using simple approximations of the complex Shi-Tomasi and Harris corner measure to create a small corner candidate set. The conventional Shi-Tomasi/Harris corner measure is then applied only to this pool of corner candidates to extract the final corners, thereby slashing the overall computation cost significantly. The proposed pruning technique is devised based on the premise that in most images, corner regions constitute a very small portion of the image. While the complex corner measure of Shi-Tomasi and Harris can be used to rank and extract the best corners, they incur excessive computations for eliminating non-corners in the entire image. This paper is organized as follows. In the next section we discuss previously reported work on accelerating Shi-Tomasi and Harris detectors. Section III provides a brief description of the Shi-Tomasi and Harris corner detectors. In Section IV, we introduce our proposed technique, named as P-Shi-Tomasi and P-Harris. Section V shows the evaluations with standard data. The paper concludes in Section VI.

II. RELATED WORK

Earlier work on corner detection has employed contour based methods that first extract the contours and then search for curvature maxima as corners. More recent corner detectors such as Shi-Tomasi and Harris are intensity based methods that compute a measure that indicates the presence of a corner directly from the intensity values and hence do not rely on a contour extraction step [7].

The Shi-Tomasi and Harris detectors are closely related in their approach to determine good corners. The corner measure is computed on every image pixel and the obvious non-corners are removed by applying a threshold on the corner measure. The corner response of pixels that are close to a good corner are also typically high and hence, a minimum distance is enforced between good corners using non-maximal suppression (NMS). Finally all the corners are sorted and only the top few corners are selected for further processing. The high computational load of the feature detection is mainly due to the complex corner measure.

In [6] and [8] the feature detection step is translated on to a Graphics Processing Unit (GPU) and the final step of NMS is parallelized. In [9], integral image is used and hence the computation time is kept constant for varying window sizes. In [10] an alternative simpler corner response is proposed that performs integer arithmetic with only multiplications and additions.

Our work proposes a technique to efficiently discard non-corners, which significantly reduces the selection and evaluation effort for the presence of corners to only corner-like regions. Unlike the work in [9] and [10], we do not compute the corner response on the entire image, hence potentially resulting in higher computation time savings. In addition, while the work in [6] and [8] achieves speedup by exploiting parallelism on the GPU, they are not well suited for low cost embedded systems. Moreover, compared to [6, 8], our technique computes the corner response only on a small set of corner candidates instead of the entire image.

III. SHI-TOMASI AND HARRIS CORNER DETECTORS

Feature detection in both Shi-Tomasi and Harris detectors is based on the local auto-correlation function that is approximated by matrix M within a small window W of each pixel $p(x,y)$:

$$M = \begin{bmatrix} \sum_W w(x) I_x^2 & \sum_W w(x) I_x I_y \\ \sum_W w(x) I_x I_y & \sum_W w(x) I_y^2 \end{bmatrix} = \begin{bmatrix} a & b \\ b & c \end{bmatrix} \quad (1)$$

Where I_x and I_y are horizontal and vertical gradients respectively and $w(x)$ is the Gaussian weight function. The eigenvalues λ_1 and λ_2 of M (where $\lambda_1 \geq \lambda_2$) indicate the type of intensity change in the window W around $p(x,y)$.

- If both λ_1 and λ_2 are small, $p(x,y)$ is a flat region.
- If λ_1 is large and λ_2 is small, $p(x,y)$ is an edge point.
- If both λ_1 and λ_2 are large, $p(x,y)$ represents a corner point.

Shi-Tomasi directly computes λ_2 as its corner response as in (2) and selects the points that have a large λ_2 as the corners.

$$\lambda_2 = \frac{(a+c) - \sqrt{(a-c)^2 + 4b^2}}{2} \quad (2)$$

Harris combines the eigen values into a single corner measure R as in (3) and avoids the explicit computation of eigen values.

$$\begin{aligned} R &= \lambda_1 \lambda_2 - k * (\lambda_1 + \lambda_2)^2 \\ &= \det(M) - k * trace^2(M) \\ &= (ac - b^2) - k * (a+c)^2 \end{aligned} \quad (3)$$

Where k is an empirical constant (k = 0.04 to 0.06).

A threshold is applied on the corner response to discard the obvious non-corners. The rest of the pixels are then ranked in the descending order of the corner response and the pixels with the highest corner response are then selected as corners after applying the non-maximal suppression.

IV. PROPOSED LOW COMPLEXITY PRUNING TECHNIQUE

We make the following observations on the Shi-Tomasi and Harris detectors:

- In most images, the obvious non-corners (i.e. the flat and edge regions) constitute a large majority of the image. Hence, the Shi-Tomasi and Harris detectors incur a lot of redundant computations as they evaluate the entire image for a high corner response.
- Expanding (2), we get

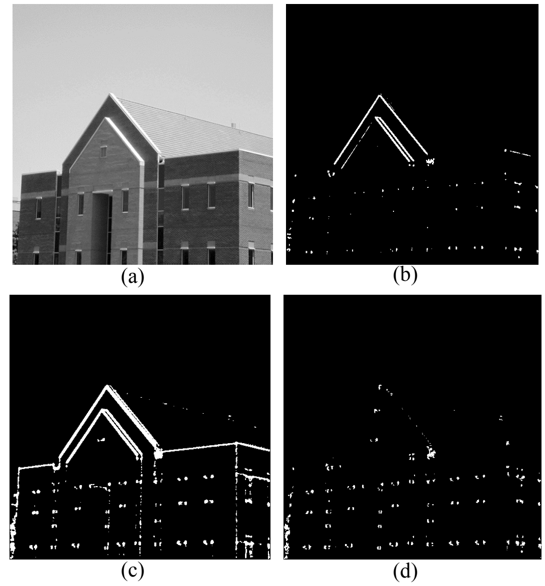

Figure 1: (a) Original image (b) Corner candidates selected using *ac* at threshold=0.05**Max*(*ac*) (c) Corner candidates selected using *a'c'* at threshold=0.05**Max*(*a'c'*) (d) Corner regions with λ_2>0.05**Max*(λ_2)

$$\lambda_2 = \frac{(a+c) - \sqrt{(a-c)^2 + 4b^2}}{2} = \frac{(a+c) - \sqrt{(a+c)^2 - 4(ac - b^2)}}{2} \quad (4)$$

  λ_2 is most influenced by the term $(ac - b^2)$ as the two $(a+c)$ terms cancel out. For a good corner, λ_2 needs to be a large value. Hence maximizing $(ac - b^2)$ which is also the *det(M)* can select good Shi-Tomasi corners without explicit eigen value computation.

- In Harris, the *trace(M)* term is introduced so that edges can also be detected. Ignoring the *trace(M)* term, the *det(M)* term alone is sufficient to select corner regions.

Based on our observations, we propose a pruning technique that approximates the *det(M),* (i.e. determinant of the auto correlation matrix M) for selecting corner candidate pixels. We then apply the conventional Shi-Tomasi or Harris corner measure on only the corner candidate set to extract the final corners.

Pixels that maximize *det(M)* also have a high value for λ_2. For a high value of $det(M) = ac - b^2$, the pixel must have a large value for ac. We propose to choose only such pixels as corner candidates. Applying an appropriate threshold can discard pixels with low ac values. Fig. 1(b) shows the corner candidates selected by applying threshold = 0.05**max(ac)*, and it is clear that this covers the final corner regions in Fig. 1(d) well.

We propose to approximate the I_x^2 and I_y^2 terms in the ac value with the absolute values for I_x and I_y as follows:

$$a' = \sum |I_x|, c' = \sum |I_y| \quad (5)$$

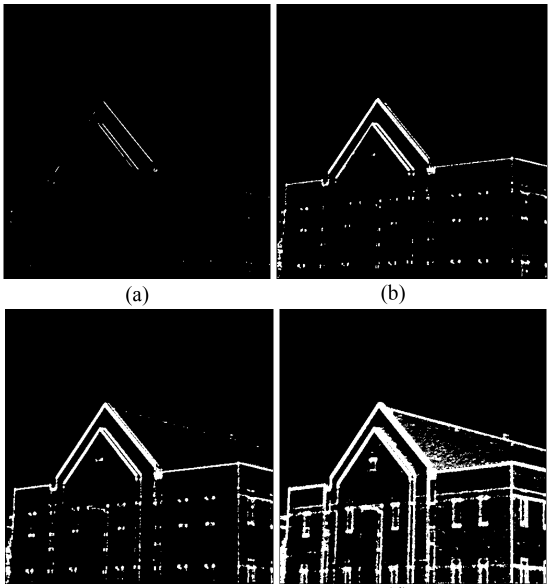

(a) (b) (c) (d)

Figure 2: *a'c'* map at various thresholds (a) 0.5 (b) 0.1 (c) 0.05 (d) 0.01

Input: Gradients I_x, I_y for Image I, feature quality threshold t

/* Pruning */

1. Compute $|I_x|$, $|I_y|$ for each pixel in I
2. Compute $a'c' = \sum|I_x| * \sum|I_y|$ for each pixel in I
3. Threshold *a'c'* image with threshold = $t * max(a'c')$ to get *Corner_Candidates*

/* Corner Response Function */

4. For each pixel in *Corner_Candidates* and its neighbors in window W
   - Compute I_x^2, I_y^2, I_xI_y
5. For each pixel in *Corner_Candidates*
   - Compute $a=\sum w(x)*I_x^2$, $c=\sum w(x)*I_y^2$, $b=\sum w(x)*I_xI_y$
   - Compute Corner Response $C = \lambda_2$ or R
   - Threshold *Corner_Candidates* with threshold = $t*max(C)$.
6. Sort in descending order of C
7. Apply non-maximal suppression

Figure 3: Outline of the P-Shi-Tomasi & P-Harris algorithm

Table 1: COMPARISON OF OPERATIONS PER PIXEL

Operations per pixel		
Methods	***Multiplications***	***Adds***
Shi-Tomasi	3 (16-bit), 3W+1 (FP), 2 (32-bit)	3W + 4 (32-bit)
Harris	4 (32 bit), 3W (FP)	3W + 3 (32-bit)
Pruning Technique	1 (16-bit)	2W (16-bit)

This eliminates the multiplication operations involved in the squared gradients. Fig. 1(c) shows *a'c'* map covers the *ac* map and the final corner regions well.

High *a'c'* value implies that *p(x,y)* has a high response for both I_x and I_y in the window W around it. Fig. 2 shows that under uniform illumination, as the threshold for *a'c'* map is reduced, corners and slanted edges are released first. This is followed by vertical and horizontal lines, and finally textures and flat regions. This shows that when the threshold to *a'c'* is applied, non-corner regions are likely to be removed, while retaining a significant amount of corners. Hence, *a'c'* can be used as an effective corner indicator measure as it elevates the corner regions above the non-corner regions. The outline of the algorithm is presented in Fig. 3.

Table 1 shows the details on the per pixel operations involved in the corner indicator used in the pruning versus the conventional Shi-Tomasi and Harris corner measures. A window size $W=3x3$ has been assumed for all the window operations. The savings in computations is achieved because the simple *a'c'* corner indicator is applied to all pixels and the more complex corner measure is applied only to the small candidate sets.

V. EVALUATIONS AND RESULTS

We use the Shi-Tomasi and Harris algorithms with a Gaussian window for computing M as in [3] with $W=3x3$ and $\sigma=0.5$ as the baseline algorithms. For the proposed technique, the pruning algorithm is first applied to the entire image in order to select the corner candidate set. Next, the corner measure of the corresponding baseline algorithm is applied to extract the final corners. We compare our P-Shi-Tomasi and P-Harris algorithms with the original Shi-Tomasi and Harris detectors in terms of the accuracy and timing. Fig. 4 shows the image dataset used, which contains four sequences ("graf", "boat", "bark" and "leuven") of images with various image transformations such changes in viewpoint, zoom, rotation and illumination [11]. For all images, we apply the thresholds on the *a'c'* and the final corner measures (λ_2 or R) such that 250, 500, 750 and 1000 corners are extracted.

A. Accuracy

The accuracy of the final corners extracted is evaluated using the repeatability rate [7], which is defined as the number of points repeated between two images with respect to the total number of detected points. Fig. 5 shows the average difference in repeatability rate between the P-Shi-Tomasi and Shi-Tomasi corners. The average difference in repeatability rate is calculated as: $\Delta r = r(proposed) - r(baseline)$. For all image transformations, the average difference in repeatability for P-Shi-Tomasi is less than 3%. It is to be noted that the difference is mostly positive and hence P-Shi-Tomasi

Figure 4: Image data set used for the evaluation (a) "graf" – viewpoint change (b) "boat" – zoom + rotation change (c) "bark" – zoom + rotation change (d) "leuven" – illumination change

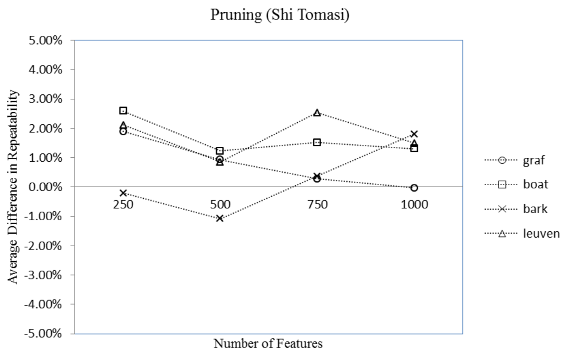

Fig. 5: Difference in average repeatability

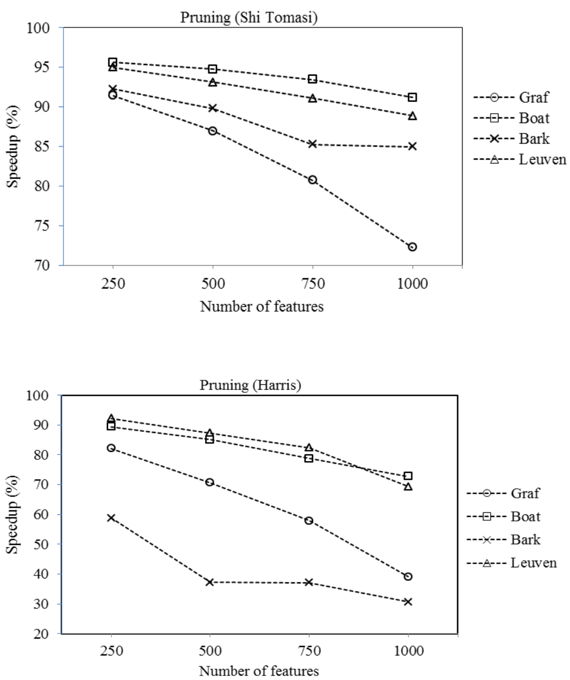

Figure 6: %Speedup in timing on NiOS-II

performs slightly better than Shi-Tomasi in most cases. The repeatability of the P-Harris is not shown as it is able to determine the same feature set as the baseline Harris detector. The results show that for the images considered, the proposed pruning method still results in the selection of final feature candidates that are exactly the same as the Harris detector and almost equivalent to the Shi-Tomasi detector.

B. *Timing*

The computational efficiency of the pruning technique is demonstrated by running the code on the NiOS-II embedded processor [12]. The timing results are obtained for the first image in the image sequence for each transformation and are presented in Fig. 6. For Shi-Tomasi, a speed up of 70-95% is observed. The average speedup of P-Harris over the Harris detector is generally lower due to the absence of the square root operation that is required by Shi-Tomasi. It can be observed that P-Harris still achieves a significant speedup of 70-90% over the Harris detector for “boat” and “leuven”. As “graf” and “bark” do not contain enough high quality features, a much lower threshold needs to be set in order to generate 1000 features. Hence larger candidate sets are evaluated, which leads to lesser speedup achieved by P-Harris. Despite this, a notable speedup of 30-80% over the Harris detector is still observed for “graf” and “bark”.

VI. Conclusions

We have presented a low cost pruning technique to accelerate the Shi-Tomasi and Harris corner detectors by using an approximate corner indicator derived from the conventional corner measure. Evaluations for repeatability showed that the corner candidates selected by the proposed pruning technique include most of the corners found by the baseline detectors. The approximate measure used for pruning allows high thresholds to be applied to remove non corner regions, while retaining a significant amount of corners. This facilitates the selection of a small but near-complete set of corner candidates, which results in significant computation savings on corner response evaluation. Experimental results demonstrate that the proposed technique achieves significant speedup in the range of 70-90% and 30-90% over Shi-Tomasi and Harris respectively for 250-1000 features. The pruning technique is well suited for high performance and low cost embedded systems. The choice of feature quality threshold to obtain the required number of features is still an open research question that is often faced in related work. As an optimized threshold results in high computation savings, our future work will focus on adaptive image thresholding techniques.

References

[1] J. Shi, C. Tomasi, “Good features to track”, Proc. of Conf. on Computer Vision and Pattern Recognition, 1994.

[2] C. Harris, M. Stephens, “A combined corner and edge detector”, Alvey Vision Conference, 1988.

[3] S. Gauglitz, T.Höllerer and M. Turk, “Evaluation of Interest Point Detectors and Feature Descriptors for Visual Tracking”, International Journal of Computer Vision, Vol 94, 2011.

[4] A. Gil, O. M.Mozos, M.Ballesta, O.Reinoso, “A comparative evaluation of interest point detectors and local descriptors for visual SLAM”, Machine Vision & Applications, Vol 21, 2010.

[5] A. Schmidt, M. Kraft, A.Kasiński, “An Evaluation of Image Feature Detectors and Descriptors for Robot Navigation”, Computer Vision and Graphics, 2010.

[6] L. Teixeira, W. Celes, M. Gattas, “Accelerated Corner Detector Algorithms”, British Machine Vision Conference, 2008

[7] C.Schmid, R. Mohr, C.Bauckhage, “Evaluation of Interest Point Detectors”, International Journal of Computer Vision, Vol 37, 2000.

[8] S. Sinha, J. Frahm, M. Pollefeys, Y. Genc. “Feature tracking and matching in video using programmable graphics hardware”, Machine Vision and Applications, 2007.

[9] P. Mainali, Qiong Yang, G. Lafruit, R. Lauwereins, L.V. Gool, “Robust Low Complexity Corner Detector”, IEEE Trans. on Circuits and Systems for Video Technology, Vol. 21, No. 4, 2011.

[10] A. Benedetti, P. Perona, “Real-Time 2-D Feature Detection on a Reconfigurable Computer,” Proc. IEEE Conf. Computer Vision and Pattern Recognition, 1998.

[11] [online] http://www.robots.ox.ac.uk/~vgg/research/affine/index.html

[12] [online] http://www.altera.com/devices/processor/nios2/ni2-index.html