Mask-based Non-Maximal Suppression with Iterative Pruning for Low-Complexity Corner Detection

Nirmala Ramakrishnan, Meiqing Wu, Siew-Kei Lam, Thambipillai Srikanthan
School of Computer Engineering
Nanyang Technological University
Singapore
{rama0056, wume0007}@e.ntu.edu.sg, {assklam, astsrikan}@ntu.edu.sg

***Abstract*—Adaptive and low-complexity corner detection is essential for the realization of modern real-time video processing applications on low-resource embedded platforms. In our earlier work, we have proposed an efficient iterative pruning strategy for widely used Shi-Tomasi and Harris corner detectors. In this paper we propose a mask-based strategy for efficient non-maximal suppression. In comparison to the conventional non-maximal suppression, applying the proposed strategy results in higher efficiency in computations for the iterative pruning based Shi-Tomasi/Harris corner detectors. On the embedded NiOS-2 platform, we achieve a speedup in execution time of 14-42% for Shi-Tomasi and 11-47% for Harris corner detectors.**

***Index Terms*—Non-maximal suppression, corner detector, low complexity, Shi-Tomasi, Harris**

I. INTRODUCTION

Modern video processing applications are increasingly being applied to low-resource embedded platforms that process video imagery that is not known a-priori. Examples include robotic vision [1] and video stabilization [2], in which feature point detection is the first step. Feature detection algorithms, Harris [3] and Shi-Tomasi [4] have been widely used in these embedded applications as they are less computationally intensive.

In our earlier work [5], we proposed a novel iterative pruning strategy for the Shi-Tomasi and Harris corner detectors that eliminated the need to manually specify a threshold on the corner measure. We were able to realize an adaptive and low-complexity corner detector by adapting the computations for corner detection to the image content.

In practice many pixels near a corner have a high corner response and non-maximal suppression is applied to select one representative pixel that is the local maximum. In [6] a simple and widely-used non-maximum suppression scheme was proposed. For each pixel considered as a potential corner candidate we scan the list of all the corners already identified and check if this pixel is in the neighborhood of any of these corners. If it is, then this pixel is suppressed; otherwise it is included as a corner.

For corner detection the user specifies the number of corners to be found and the minimum distance that needs to be enforced with non-maximal suppression between the corners as these parameter values are application dependent. The classic non-maximal suppression scheme works well when: (a) the required number of corners to be found is small and/or (b) the minimum distance between corners is small. However when either of these parameter values is increased – in order to get a dense corner set or a well-spread out corner set, the computations for non-maximal suppression also increase. This is because many more corner candidates need to be processed before the required number of corner candidates is found.

Pseudocode for Conventional Non-Maximal Suppression

```
For each corner candidate p
  For each already selected corner k
     If distance(p,k) < min_distance
        Discard p;
        Break; //Go to next p
     Endif
  EndFor
  Add p to the list of selected corners
EndFor
```

Fig.1 Conventional Non-Maximal Suppression

As we need high compute efficiency in corner detection, it is important to overcome the computational complexity for the non-maximal suppression step. In [7] a parallel NMS algorithm was proposed for GPU implementation.

In this paper we propose to use a mask-based NMS strategy similar to [7] that works well with the iterative pruning based corner detectors [5]. As corners are detected in several iterations, the proposed strategy allows the suppression of neighborhood pixels in subsequent iterations based on the corners detected in the current iteration. When a corner candidate pixel is suppressed in the iterative pruning based corner detection, this pixel does not undergo the complex corner measure computation thereby increasing the computation efficiency of such detectors.

II. PROPOSED MASK-BASED NON-MAXIMAL SUPPRESSION STRATEGY

Shi-Tomasi and Harris corner detectors first compute a corner measure for each pixel in the image. Generally, the corner response is high for many pixels in the vicinity of a corner. In order to determine a single representative pixel non-maximal suppression is applied such that the pixel that is the

978-1-4799-4833-8/14/$31.00

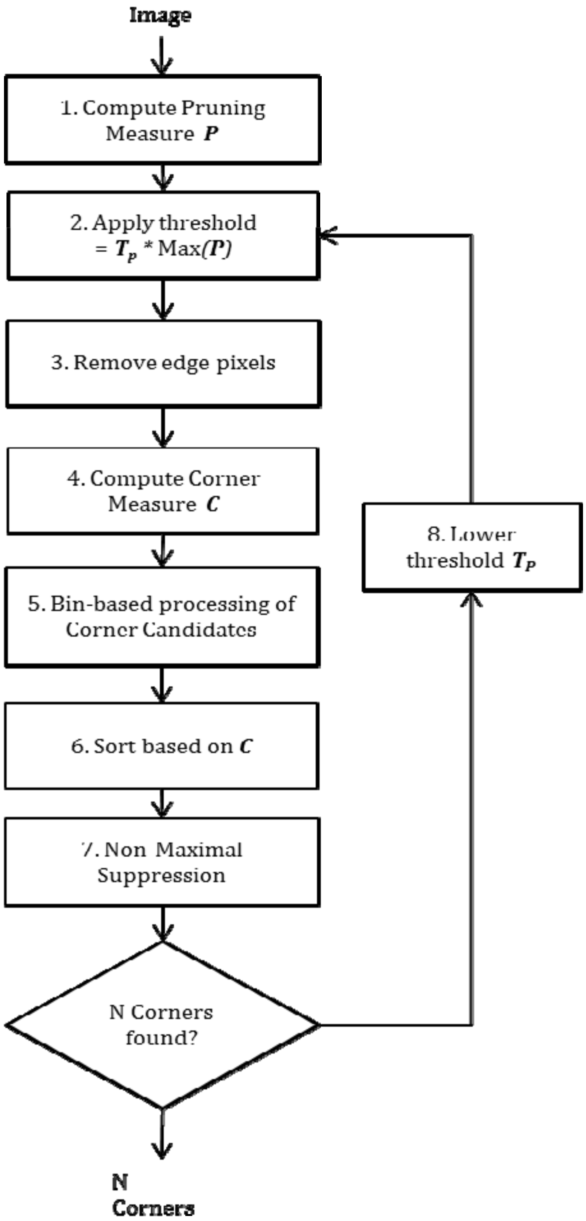

Fig. 2 Iterative Pruning based Corner Detection

local maximum in the corner measure in a small neighborhood is selected. Fig. 1 shows the classical approach to non-maximal suppression in corner detection as given in [6]. Note that for each candidate we check if it is in the neighborhood of any of the already selected corner pixels. This size of the neighborhood is defined by the minimum distance to be enforced between selected corners.

The iterative pruning based corner detectors [5] select corners in several iterations. Fig. 2 shows the flow for this iterative processing. A low-complexity pruning measure P is computed for all pixels in the image. As the pruning measure is derived from the corner measures C for Shi-Tomasi and Harris, pixels with high values for the pruning measure are also likely to have a high value for the corner measure. Therefore, we release corner candidates by applying a threshold based on the pruning measure P (Step 2). An edge response removal step (Step 3) is used to further eliminate pixels likely to be edges and results in a small corner candidate pool. The corner measure C is finally computed only on this small pool of corner candidates. These candidates are then crudely sorted into corner measure bins depending on their corner measure values C. As shown in [5] bins with high quality corner candidates do not grow in size even when more candidates are released in subsequent iterations of lowering the threshold on the pruning measure P. These candidates are then collected, sorted based on their corner measure C and finally the non-maximal suppression is applied to select a single representative pixel in a corner region. If the required number of corners is found at the end of this step, the algorithm stops – if not, additional iterations of releasing corner candidates is invoked by lowering the threshold on the pruning measure P (Step 8).

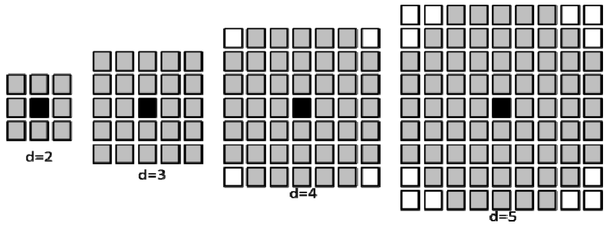

Fig. 3. Non-maximal suppression masks for various minimum distance values, d. Grey pixels are suppressed by the black pixel.

The conventional non-maximal suppression in [6] only checks if the current corner candidate is in the vicinity of already selected corners. However, in the iterative pruning based corner detectors, once a batch of corners is selected in the current iteration, and another iteration is needed for selecting more corners, it is possible to filter out future corner candidates in the neighborhood of the already selected corners from being processed. This will avoid costly corner measure computations on these candidates.

To achieve this, we propose a mask-based non-maximal suppression strategy that maintains the status of every pixel in the image as either SUPPRESSED or NOT_SUPPRESSED. This map of status flags is the same size as the image and is initialized to NOT_SUPPRESSED. For any pixel to be chosen as a corner candidate – it needs to be NOT_SUPPRESSED. When a pixel is chosen as a corner in Step 7, we update the neighborhood around this pixel to all SUPPRESSED. The size of the neighborhood is determined by the minimum distance parameter and it can be represented as a pre-computed circular mask as shown in Fig. 3. At Step 2, when releasing new corner candidates in a subsequent iteration, we check if the candidate is also NOT_SUPPRESSED. If not, we discard this corner candidate from further costly processing for corner measure. In effect, we only compute the costly corner measure C for corner candidates which are not in the vicinity of already selected

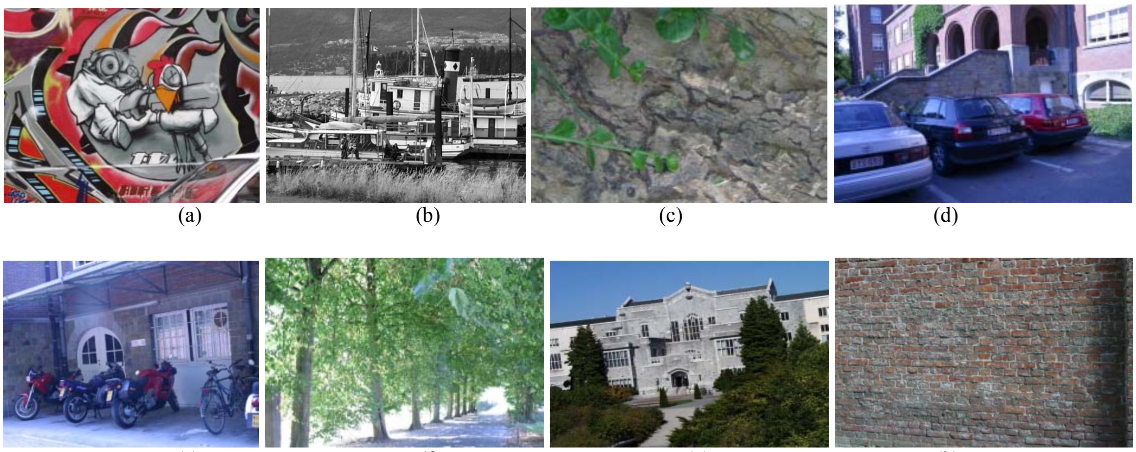

(a) (b) (c) (d)

(e) (f) (g) (h)

Fig. 4. Image dataset used for the evaluation (a) "graf" (800x640) (b) "boat" (850x680) (c) "bark" (765x512) (d) "leuven" (900x600) (e) "bike" (1000x700) (f) "trees" (1000x700) (g) "ubc" (800x640) (h) "wall" (1000x700)

corners and improve the overall compute-efficiency of the iterative pruning based corner detectors.

TABLE I

REDUCTION IN THE NUMBER OF CORNERS FOR SHI-TOMASI

	N=300			*N*=1000		
Images	*d*=5	*d*=10	*d*=15	*d*=5	*d*=10	*d*=15
graf	7%	16%	34%	7%	19%	32%
boat	2%	10%	17%	5%	24%	32%
bike	2%	13%	28%	9%	28%	39%
bark	2%	10%	24%	4%	16%	26%
leuven	5%	14%	35%	9%	27%	39%
ubc	5%	12%	54%	38%	41%	51%
trees	2%	11%	21%	3%	11%	35%
wall	2%	8%	15%	2%	20%	40%

TABLE II

REDUCTION IN THE NUMBER OF CORNERS FOR HARRIS

	N=300			*N*=1000		
Images	*d*=5	*d*=10	*d*=15	*d*=5	*d*=10	*d*=15
graf	7%	21%	47%	10%	26%	41%
boat	33%	39%	30%	37%	46%	53%
bike	9%	21%	37%	16%	38%	52%
bark	5%	21%	32%	30%	37%	50%
leuven	6%	45%	60%	16%	44%	58%
ubc	6%	19%	36%	36%	50%	59%
trees	4%	9%	16%	40%	47%	50%
wall	1%	42%	48%	37%	25%	57%

III. EVALUATIONS

We use the images in Fig. 4 for our evaluations. We compare the performance of the proposed mask-based NMS strategy with the conventional NMS strategy as in [6] when both are applied to the iterative pruning based Shi-Tomasi and Harris corner detectors. We show the results with varying number of corners to be detected, specified as N and the minimum distance to be enforced by non-maximal suppression, specified as d. The low-complexity pruning applied in steps 1-3 in the iterative pruning strategy as shown in Fig. 2 results in a small pool of corner candidates that undergo the complex corner measure computation. Tables I and II show the reduction in this number of corner candidates when the mask-based NMS strategy replaces the conventional NMS strategy for the iterative pruning based detectors. Clearly, when the minimum distance d between the corners is higher, the overall reduction in the number of candidates also increases. A larger value for d implies that more number of corner candidates need to be processed to get the required number of N corners. In this scenario the proposed mask-based NMS strategy is able to avoid the processing of more number of candidates that are already in the vicinity of corners chosen in earlier iterations. Similarly when the number of corners to be detected, N, is increased, for many images, this results in higher reduction of the number of corner candidates when the proposed mask-based NMS strategy is applied. The exception to this behavior is when in some images rich in corners, such as “boat” the corner candidates released in the final iteration may be large and sufficient to detect the required number of corners for all the values of minimum distance d and/or N considered in our evaluations. In such cases, we see that the reduction in number of candidates does not increase when we consider a higher value of d.

We demonstrate the improvement in performance due to the savings in computations we achieve by running the algorithms on the NiOS-2 embedded platform [8].

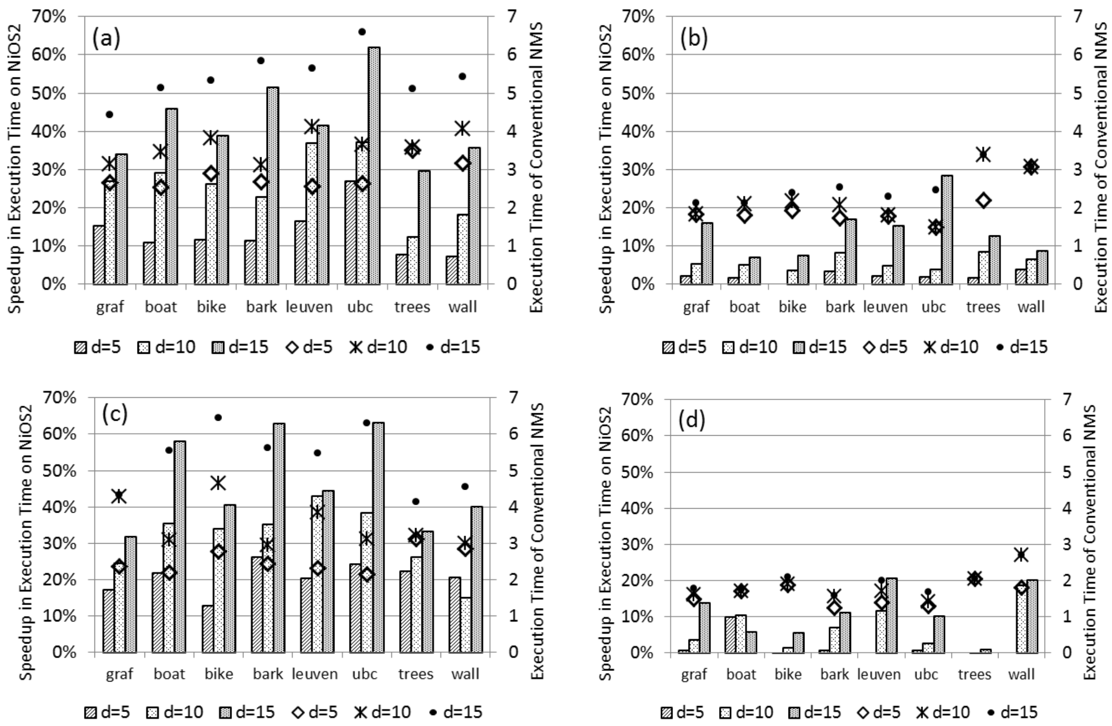

Fig. 5. Evaluation results: Speedup achieved by the proposed mask-based NMS strategy compared to the conventional NMS in the iterative pruning based Shi-Tomasi/Harris corner detectors (a) Shi-Tomasi N=1000 (b) Shi-Tomasi N=300 (c) Harris N=1000 (d) Harris N=300

We report the results as the execution time of the conventional NMS strategy and the relative speedup in execution time achieved with the proposed mask-based NMS strategy in Fig. 5. In general, the speedup is higher when we increase the minimum distance *d* for non-maximal suppression or the number of corners to be detected *N*, similar to our analysis in Tables I and II. For a minimum distance of 15 pixels, a speedup of 14% and 42% is achieved with 300 and 1000 Shi-Tomasi corners respectively. For Harris, a speedup of 11% and 47% is achieved for 300 and 1000 corners respectively.

IV. Conclusions

Non-maximal suppression is a common step in all corner detectors. In this paper we have proposed a mask-based NMS strategy which can be combined with iterative pruning based Shi-Tomasi/Harris corner detectors, from our earlier work. The proposed strategy exploits the characteristic of corner selections in several iterations as performed by the iterative pruning based corner detectors and achieves savings in computations by suppressing corner candidates before the complex corner measure is applied to them. Our evaluations on an embedded processor platform show that the proposed mask-based NMS strategy achieves savings in computations especially when the minimum distance for non-maximal suppression and/or the number of corners to be detected is high. Therefore it is well suited for the iterative pruning based Shi-Tomasi and Harris corner detectors that achieve low-complexity and adaptive corner detection.

References

[1] K. Schmid and H. Hirschmuller, "Stereo vision and IMU based real-time ego-motion and depth image computation on a handheld device," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, 2013, pp. 4671-4678.

[2] T. Tsung-Han, F. Chih-Lun, and C. Hui-Min, "Design and Implementation of Efficient Video Stabilization Engine Using Maximum a Posteriori Estimation and Motion Energy Smoothing Approach," *Circuits and Systems for Video Technology, IEEE Transactions on,* vol. 22, pp. 817-830, 2012.

[3] C. Harris and M. Stephens, "A combined corner and edge detector," in *Proc. Fourth Alvey Vision Conference*, 1988, pp. 147-151.

[4] J. Shi and C. Tomasi, "Good features to track," in *Computer Vision and Pattern Recognition, 1994. Proceedings CVPR '94., 1994 IEEE Computer Society Conference on*, 1994, pp. 593-600.

[5] N. Ramakrishnan, W. Meiqing, L. Siew-Kei, and T. Srikanthan, "Automated thresholding for low-complexity corner detection," in *Adaptive Hardware and Systems (AHS), 2014 NASA/ESA Conference on*, 2014, pp. 97-103.

[6] C. Tomasi and T. Kanade, "Detection and Tracking of Point Features," Carnegie Mellon University 1991.

[7] L. Teixeira, W. Celes, and M. Gattass, "Accelerated Corner-Detector Algorithms," *British Machine Vision Conference,* 2008.

[8] *Nios II Processor: The World's Most Versatile Embedded Processor*. Available: http://www.altera.com/devices/processor/nios2/ni2-index.html