# Accelerated Embedded AKAZE Feature Detection Algorithm on FPGA

Lester Kalms Ruhr-Universität Bochum (RUB) Bochum, Germany lester.kalms@rub.de Khaled Mohamed German University in Cairo (GUC) Cairo, Egypt khaled.essam@student.guc.edu.eg Diana Göhringer Technische Universität Dresden Dresden, Germany diana.goehringer@tu-dresden.de

# ABSTRACT

Feature detection is a major operation in various computer vision systems. The KAZE algorithm and its improved version, Accelerated-KAZE (AKAZE), are considered as the first algorithms to detect features by building a scale space using nonlinear diffusion. However, the detection part of the algorithm achieves only 6 fps on an Intel Core i7-4790 processor for an image resolution of 1024x768. This work proposes a pipelined architecture for a hardware accelerator that performs the AKAZE detection algorithm. Firstly, modifications are done to the original algorithm to reduce the amount of computations and memory accesses. Then, the accelerator is implemented on a Xilinx Zynq SoC and achieves 98 fps for the same resolution and a frequency of 100 MHz. Compared to the original algorithm, the design has an error in average inliers ratio by only 4%.

## **CCS CONCEPTS**

• Computing methodologies  $\rightarrow$  Computer vision; Object detection; • Hardware  $\rightarrow$  Reconfigurable logic and FPGAs; Hardware accelerators;

## **KEYWORDS**

feature detection, sliding window, FPGAs, Xilinx Zynq SoC, fast retina keypoint, embedded application, image processing

#### **ACM Reference Format:**

Lester Kalms, Khaled Mohamed, and Diana Göhringer. 2017. Accelerated Embedded AKAZE Feature Detection Algorithm on FPGA. In *Proceedings of HEART2017, Bochum, Germany, June 07-09, 2017,* 6 pages. https://doi.org/10.1145/3120895.3120898

# **1 INTRODUCTION AND MOTIVATION**

Detecting features in an image is considered as a fundamental step in many computer vision applications such as object recognition, image registration and motion tracking [14]. There are two basic criteria for choosing a detection algorithm. The first criterion is the ability of the algorithm to be used in widely varying applications, while the second one is the repeatability of the algorithm [14]. Repeatability is defined as the ability of detecting the same feature

HEART2017, June 07-09, 2017, Bochum, Germany

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5316-8/17/06...\$15.00

https://doi.org/10.1145/3120895.3120898

in different locations invariantly to rotation, scale, and change in brightness or contrast.

KAZE and accelerated KAZE (AKAZE), introduced in [4] and [16] respectively, are the first algorithms to use the non-linear diffusion in multi-scale feature detection. Other approaches use linear diffusion to create the scale-space. According to Alcantarilla et al. [16], both KAZE and AKAZE show better repeatability and performance than other algorithms such as ORB [19], BRISK [12], SIFT [13] and SURF [5].

The main problem with AKAZE is the computation intensiveness. For an image of resolution of 1024x768, we measure that the detection part of the original algorithm achieves only 6 fps in average on an Intel Core i7-4790 processor. In order to achieve higher frame rates and meet real-time requirements, this paper introduces a hardware acceleration of the detection part of the AKAZE algorithm. Additionally, the detected features (or keypoints) are sent to the FREAK descriptor [3], which is processed in software and used instead of the MLDB descriptor of AKAZE. The FREAK descriptor is faster to compute and has a much lower memory load.

The paper is organized as follows: Section 3 shows a brief description of the AKAZE detection algorithm. Section 4 introduces the modifications done on the original algorithm. A detailed hardware implementation is described in section 5. In Section 6, the performance of the design and the hardware utilization are represented. Section 7 concludes this work and gives an outlook.

# 2 STATE OF THE ART

Various algorithms are introduced to detect image features in multiscale framework. SIFT [13] and SURF [5] are the most widely used detectors in computer vision. Both are using the Gaussian (linear) diffusion to build the scale-space. Other algorithms, such as FAST [18], ORB [19] and BRISK [12], are proposed to reduce the cost of heavy computations. Meanwhile, all these algorithms, in addition to AKAZE, are not applicable for high frame rate especially for high resolution images. Hardware implementations are designed to accelerate those algorithms. A throughput of 30 fps is obtained using SIFT and SURF for VGA resolution (640x480) [7-9, 15]. An architecture based on ORB is designed in [11] producing 55 fps for VGA resolution. The highest throughput achieved is 127 fps for full HD images using optimized AKAZE algorithm [10]. One missing part in this paper is the computation of the contrast factor, which is essential in building the non-linear scale-space. Furthermore, the hardware testing methodology of this paper is very vague.

## **3 AKAZE FEATURE DETECTION**

The AKAZE detection algorithm consists of 3 parts: computing the contrast factor, building the non-linear scale-space and detecting

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HEART2017, June 07-09, 2017, Bochum, Germany



Figure 1: Block diagram of the complete system

features. They are shown in the lower part of Figure 1. This section gives a summarized introduction to the different parts.

## 3.1 Computing the Contrast Factor

The contrast factor is significantly important in building the nonlinear scale-space. Firstly, the image is smoothed by a Gaussian filter. Second step is to calculate the maximum absolute gradient value of the smoothed image, called hmax. This is done by looping on the whole image, calculating the absolute value of the gradient of each pixel. Third step is to fill a histogram of 300 *bins* with the gradient values divided by the parameter  $h_{\text{max}}$ . The following step is to loop on the histogram to find the histogram index *i* at which the 70% percentile of the gradient histogram is achieved. The contrast factor can then be calculated by:

$$k = \frac{h_{\max} \cdot i}{300} \tag{1}$$

## 3.2 Building the Nonlinear Scale-Space

The scale-space levels are built by numerically solving the partial differential equation (Eq. 2) iteratively using the Fast Explicit Diffusion scheme (FED) with varying time steps  $\tau_i$  (Eq. 3).

$$\frac{\partial L}{\partial t} = div(g(|\nabla L_{\sigma}|) \cdot \nabla L) \tag{2}$$

$$L^{j+1} = (I + \tau_j A(L^j))L^j$$
(3)

The function  $g(|\nabla L_{\sigma}|)$  is called the conductivity function (or flow function) and A(L) is the matrix notation of this function. The conductivity function is chosen to be:

$$g(|\nabla_{\sigma}|) = \frac{1}{1 + \frac{|\nabla L_{\sigma}|^2}{k^2}} = \frac{k^2}{k^2 + |\nabla L_{\sigma}|^2}$$
(4)

where k is the contrast factor. To prevent using two division operations, Eq. 4 has been adjusted. For each pixel, computing the conductivity function (Eq. 4) requires 1 division, 2 multiplication and 2 addition operations. Additionally, each FED step (Eq. 3) needs 5 multiplication and 12 add/subtract operations for each pixel. That is what makes AKAZE computationally intensive and, consequently, Lester Kalms, Khaled Mohamed, and Diana Göhringer



Figure 2: Scale-space representation in AKAZE (3 octaves, 4 sub-levels)

time consuming. The scale-space in AKAZE (see Figure 2) is a pyramidal framework. It consists of octaves and each octave consists of sub-levels. Each octave is quarter the size of the previous.

#### 3.3 Feature Detector

AKAZE uses the determinant of the Hessian (DoH) blob-detector. After building the non-linear scale-space, DoH images are computed at increasing scale sizes as the sub-levels increases. That is why AKAZE down-samples the images after certain sub-levels.

The keypoints (or features) are extracted by comparing the pixels in DoH images with a surrounding window of size 3x3. If the value of this pixel is greater than the other 8 surrounding pixels and a predefined threshold, the pixel is compared spatially with keypoints in the same sub-level in a window of radius  $\sigma^2$  to exclude the repeated keypoints that exist inside the same circle. The same comparison is made with the upper and lower level and the radius depends on the sub-level.

## 4 MODIFIED DESIGN

Four major modifications are made to the original implementation to reduce computations, memory access and array comparisons.

## 4.1 Fixed Point Implementation

The original algorithm is implemented in software using 32-bit floating point numbers. The implementation of floating point calculations is very costly in terms of hardware resources. So, the whole implementation of AKAZE is retyped to fixed point representation. To keep the same accuracy, all produced images, such as derivative, flow and smoothed images, are in 16-bit format except for the DoH images that are in 32-bit format.

## 4.2 Computing the Contrast Factor

The parameter  $h_{\text{max}}$  is used two times: the first time is when normalizing the gradient values to be fit into 300 *bin* histogram and another time to be multiplied by the ratio *i*/300 as in (Eq. 1). From measurements with different image and random input data, it is found that  $h_{\text{max}}$  does not exceed the value 0.5. The original implementation requires looping the whole image twice, one to find  $h_{\text{max}}$ and another to fill the histogram. Instead, in the modified design the image is looped only once. This can be done by normalizing the gradient values to 0.5 instead of  $h_{\text{max}}$  The histogram length is not more than 300 *bins* and can take any arbitrary length. So, a new parameter *nbin*<sub>max</sub>, which indicates the maximum length that is achieved, is taken into consideration. Now, only one loop is needed to fill the histogram with the new normalized gradient values and to calculate *nbin*<sub>max</sub> and  $h_{\text{max}}$  for the next step. The next step is to compute the contrast factor by (Eq. 5).

$$k = \frac{h_{max} \cdot i}{nbin_{max}} \tag{5}$$

This new method doubles the speed of computing the contrast factor and reduces the memory access by half. Another change is to approximate the square root operation, for computing the gradient, into one division, one multiplication and one summation, as in (Eq. 6), where  $L_{max}$  and  $L_{min}$  are the maximum and minimum values of the first derivatives  $L_x$  and  $L_y$  respectively.

$$\nabla L \sqrt{L_x^2 + L_y^2} \approx |L_{max}| + \frac{1}{2} \frac{L_{min} \cdot L_{min}}{|L_{max}|} \tag{6}$$

Additionally, the contrast factor is redefined to be 68.75% percentile instead of 70%. This replaces multiplication by 0.7 to be multiplication by 0.6875, or in other words multiplication by (1/2 + 1/8 + 1/16), which can be achieved by adding and shifting instead of multiplying.

# 4.3 Keypoint Comparison

Each detected keypoint in AKAZE is compared to all other keypoints in the keypoint buffer, to avoid replication of neighboring keypoints. This method has a complexity of  $O(n^2)$  for n keypoints, which makes it very computationally intensive. The modified design takes the advantage that the keypoints are detected in ascending order with respect to the vertical position. So, adaptive searching pointers are used to determine the beginning and ending of the keypoints to be compared in the buffer. This adaptive searching decreases the complexity dramatically from  $O(n^2)$  to be approximately  $O(n \cdot log(n))$ . If the keypoints are only compared to the lower and the same level, the number of matches differs by about 1%. So, comparing with the upper level is excluded.

#### 4.4 FREAK Descriptor

Fast Retina Keypoint (FREAK) is a descriptor that is inspired by the human visual system and more precisely the retina [3]. In our experiments, we have compared the MLDB descriptor of AKAZE with other descriptors, like BRIEF [6], BRISK or FREAK. The FREAK descriptor is in general faster to compute and has a lower memory load and is more robust than the other algorithms. The greatest advantage over the MLDB descriptor is the low memory load. Since the MLDB descriptor needs the computed images and derivatives of all levels in scale-space to be stored into memory, which is critical in an embedded system with limited memory resources and bandwidth. Our software tests have shown that the FREAK descriptor achieves comparable or better results than other descriptors in terms of correct matches for AKAZE.



Figure 3: Block diagram of 5x5 window generator



Figure 4: Block diagram of 5x5 window generator (replicated borders)

## **5 HARDWARE IMPLEMENTATION**

This section describes different hardware blocks used to build the AKAZE accelerator on an FPGA, in addition to the software used to test the accelerator.

#### 5.1 System Overview

Row Buffe

\_Image Width - Horizontal Window Size\_

The system consists of three parts (see Figure 1) PC platform, ARM processor and the hardware accelerator. Both ARM and accelerator are embedded into the Zynq-7000 chip on the ZedBoard [2]. The PC platform contains the software that reads the images, sends them to the ARM via Ethernet and then receives the keypoints. It also contains the description and matching parts that use these keypoints to test the system.

The accelerator has two pipelined stages: stage 1 computes the contrast factor and stage 2 builds the non-linear scale space and detects the keypoints. The ARM sends an image (first frame) to stage 1, and when it finishes, it sends an interrupt to the ARM to send the same image to stage 2. The contrast factor of an image is computed only once for a frame. Stage 2 has two outputs that are saved in memory: the keypoints of the first octave and the resized image for the following octave. When stage 2 finishes, it sends an interrupt to the ARM to send the resized image back from the memory to stage 2 to detect the keypoints for the second octave. This process is repeated until all keypoints of different octaves are sent to the ARM. While stage 2 is working on the first frame, stage 1 computes the contrast factor for the second frame. So, both stages are working simultaneously but for different frames.

#### HEART2017, June 07-09, 2017, Bochum, Germany

Pixels Buffe

\_Horizontal Window Size

Pixels



Figure 5: Block diagram of stage 1 of the accelerator

## 5.2 Window Generator

Most modules in the implemented hardware design must scan the whole image and use a specific number of pixels as a window to generate its result. The implemented modules use different window sizes ranging from 3x3 to 17x17. An extended architecture of the window generator, like the one proposed in [17], is used. The window generator (see Figure 3) consists of pixels buffer (on the right) of shift registers and row buffer (on the left), built with BRAM. The output of the window generator is the content of the pixels buffer. In the extended design (see Figure 4), some multiplexers are added to the pixels buffer to support horizontal and vertical replicated borders. The extended design also supports variable image widths to be used by different octaves.

## 5.3 Mathematical Operations

In the AKAZE hardware design, there are five mathematical operations that use the output of the window generator: Gaussian, Scharr derivative, conductivity function, FED and determinant of Hessian. Each function of them uses a different size of the window generator according to the kernel size and the scale.

# 5.4 Contrast Factor

A block diagram of computing the contrast factor is shown in Figure 5. The histogram is a simple dual port BRAM (A for writing and B for reading) of length 1024. The histogram controller is a Finite State Machine (FSM) that has three states: filling histogram, reading histogram and halt. For the first state, the controller enables writing and sets both addresses A and B to the index value *nbin* so that the value at address *nbin* is incremented by one. In the second state, the controller reads the histogram values (through address B) starting from the first index until 68.75% percentile of the gradient histogram is achieved. Then, the contrast factor is calculated by (Eq. 5), where i = AddrB.



Figure 6: Non-linear scale-space generation

## 5.5 Non-Linear Scale-Space

The hardware is designed for building 4 sub-levels per octave. Each sub-level is generated by a "NL Scale-Space" block. This block consists of (Figure 6): Gaussian Filter, conductivity function generator and a number of FED steps. Because the pixel stream of the flow image (generated by the conductivity filter) and the original image must be presented simultaneously, a FIFO is needed before the first FED step to compensate the delay of the window generator of the Gaussian and conductivity filters. The number of FED steps varies according to the sub-level and octave. In sub-level 0 for octave 0, there are no FED steps. So, this sub-level is built by only a Gaussian filter with window size 7x7 to remove the noise of the original image. Therefore, multiplexers are used in octave 0 to choose between this Gaussian filter and the normal "NL Scale-Space step 0" block.

## 5.6 Keypoint Detector

The DoH filter is applied on each input image with a different scale, which depends on the sub-level (Figure 7). Taking the advantage of the linearity property of the convolution process, the second order derivatives can be computed using the second order derivative kernel without computing the first order derivative images. The second order derivative kernel is calculated by convoluting the two derivative kernels. If more than one keypoint is present at the same time from different "local extrema" blocks, the multiplexer chooses the keypoint with the lower sub-level first. New keypoints are compared to the keypoint stored in the buffer by the adaptive search pointers. The keypoint is a 64-bit vector, which contains the x and y position, the sub-level number and the response value

#### Accelerated Embedded AKAZE Feature Detection Algorithm on FPGA



Figure 7: Block diagram of the keypoints detector

(DoH value). According to the measurements, the buffer length is set to a maximum of 8192 keypoints.

# 6 **RESULTS**

#### 6.1 Detection Repeatability

A set of 48 images is used to test the repeatability of the system. The images are in grey-scale and differs in blurring, viewpoint, rotation, brightness, JPEG compression and noise added. They are taken from the Oxford Visual Geometry Group datasets [1] and the "iguazu" dataset presented by Alcantarilla et al. [16]. The original implementation of Alcantarilla et al. [16] has been taken for comparison and compiled using the O2 optimization flag. The different configuration flags only differ in the contrast factor as described in section 4.2.

Figure 8 shows a comparison between the original algorithm and the modified one for the mentioned datasets with respect to the inliers ratio. That is, the ratio between correct matches to total matches. Generally, our modified algorithm shows the same behavior as the original one for different images. The original algorithm gives 76% average inliers ratio and ours gives 73% for the test datasets. Consequently, an error of only 4% in the average inliers ratio is present. The reduced ratio in some datasets is mainly due to the omitted sub-pixel refinement when using the FREAK descriptor, which can be added without increasing the computation time.

#### 6.2 System Throughput and Resources

The hardware is synthesized and implemented using Xilinx Vivado 2016.4 and the ZedBoard [5] as the evaluation and development board. The implemented design gives a maximum operating frequency of 100 MHz. The throughput of the design depends on the size of the input frame. Eq. 7 is used to calculate the time of the system:

$$T_{tot} = \frac{4}{3}W(H+n)(1-4^{-o})/freq$$
(7)

Where W is the image width, H is the image height, O is the number of octaves and n is a factor representing the initial delay required to fill the row buffers in the window generators. The parameter n is estimated to be 4 for stage 1 and 25 for stage 2 (Figure 1). By measuring the computation time of the accelerator using the ARM processor, it is found that (Eq. 7) is completely accurate. In our system, we are using an image resolution of 1024x768 and 2 octaves. Stage 2 determines the overall computation time, since O



Figure 8: Inliers ratio for different image sets

is only 1 in stage 1. So, the accelerator needs a total time of 10.15 ms or 98 fps.

Table 1 shows resource utilization for basic building blocks and the overall system as given by the implementation tool. The results show that it is possible to fit AKAZE feature detection algorithm on a low-cost FPGA board for embedded systems, with remaining resources for light weighted a descriptor like FREAK.

#### 7 CONCLUSION AND FUTURE WORK

This paper introduces a real-time hardware implementation for the AKAZE detection algorithm. The design consists of two parallel stages embedded on an FPGA connected with an ARM processor. The system is evaluated and tested on a ZedBoard and runs on a frequency of 100 MHz. For a resolution of 1024x768, the system gives a throughput of 98 fps while keeping the same accuracy as the original algorithm implemented on software. In future work, the FREAK descriptor can be implemented in hardware and integrated in parallel to the AKAZE accelerator as a third stage. Moreover, using the sub-pixel refinement to reallocate the keypoints shows

HEART2017, June 07-09, 2017, Bochum, Germany

Table 1: System utilization

Block	FF	LUT	BRAM	DSP
Gaussian 5x5	626	720	2	7
Gaussian 7x7	593	843	3	10
Conductivity	1248	966	1	2
FED	641	327	2	5
DoH Level 0	1829	1452	4	2
DoH Level 1	3279	1601	6	2
DoH Level 2	3279	1598	6	2
DoH Level 3	5249	1942	8	2
Local Extrema	352	365	2	0
Resize Image	175	187	0.5	0
Contrast Factor	2932	2704	4.5	13
NL Scale-Space	19806	13803	55.5	136
Keypoints Detection	15576	8438	48	8
Total System	38314	24945	108	157
	(36%)	(47%)	(77%)	(71%)

better repeatability results. The sub-pixel refinement uses the DoH images, so this can be implemented in parallel to the local extrema block.

# ACKNOWLEDGMENTS

This work is part of the TULIPP project, which is funded by the European Commission under the H2020 Framework Programme for Research and Innovation under grant agreement No 688403.

#### REFERENCES

- 2004. Visual Geometry Group Home Page. (2004). Retrieved February 20, 2017 from http://www.robots.ox.ac.uk/~vgg/data/data-aff.html
- [2] 2014. ZedBoard HW Users Guide. (2014). Retrieved February 20, 2017 from http:// zedboard.org/sites/default/files/documentations/ZedBoard\_HW\_UG\_v2\_2.pdf
- [3] A. Alahi, R. Ortiz, and P. Vandergheynst. 2012. FREAK: Fast Retina Keypoint. In 2012 IEEE Conference on Computer Vision and Pattern Recognition. 510–517. https://doi.org/10.1109/CVPR.2012.6247715
- [4] Pablo Fernández Alcantarilla, Adrien Bartoli, and Andrew J. Davison. 2012. KAZE Features. Springer Berlin Heidelberg, Berlin, Heidelberg, 214–227. https://doi. org/10.1007/978-3-642-33783-3\_16
- [5] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. 2006. SURF: Speeded Up Robust Features. Springer Berlin Heidelberg, Berlin, Heidelberg, 404–417. https: //doi.org/10.1007/11744023\_32
- [6] Michael Calonder, Vincent Lepetit, Christoph Strecha, and Pascal Fua. 2010. BRIEF: Binary Robust Independent Elementary Features. Springer Berlin Heidelberg, Berlin, Heidelberg, 778–792. https://doi.org/10.1007/978-3-642-15561-1\_56
- [7] L. C. Chiu, T. S. Chang, J. Y. Chen, and N. Y. C. Chang. 2013. Fast SIFT Design for Real-Time Visual Feature Extraction. *IEEE Transactions on Image Processing* 22, 8 (Aug 2013), 3158–3167. https://doi.org/10.1109/TIP.2013.2259841
- [8] F. C. Huang, S. Y. Huang, J. W. Ker, and Y. C. Chen. 2012. High-Performance SIFT Hardware Accelerator for Real-Time Image Feature Extraction. *IEEE Transactions* on Circuits and Systems for Video Technology 22, 3 (March 2012), 340–351. https: //doi.org/10.1109/TCSVT.2011.2162760
- [9] D. Jeon, M. B. Henry, Y. Kim, I. Lee, Z. Zhang, D. Blaauw, and D. Sylvester. 2014. An Energy Efficient Full-Frame Feature Extraction Accelerator With Shift-Latch FIFO in 28 nm CMOS. *IEEE Journal of Solid-State Circuits* 49, 5 (May 2014), 1271–1284. https://doi.org/10.1109/JSSC.2014.2309692
- [10] Guangli Jiang, Leibo Liu, Wenping Zhu, Shouyi Yin, and Shaojun Wei. 2015. A 127 Fps in Full Hd Accelerator Based on Optimized AKAZE with Efficiency and Effectiveness for Image Feature Extraction. In Proceedings of the 52Nd Annual Design Automation Conference (DAC '15). ACM, New York, NY, USA, Article 87, 6 pages. https://doi.org/10.1145/2744769.2744772
- [11] K. J. Lee K. Y., Byun. 2013. A hardware design of optimized orb algorithm with reduced hardware cost. Advanced Science and Technology Letters 43 (Dec 2013), 58–62. https://doi.org/10.1023/B:VISI.0000029664.99615.94

- [12] S. Leutenegger, M. Chli, and R. Y. Siegwart. 2011. BRISK: Binary Robust invariant scalable keypoints. In 2011 International Conference on Computer Vision. 2548– 2555. https://doi.org/10.1109/ICCV.2011.6126542
- [13] David G. Lowe. 2004. Distinctive Image Features from Scale-Invariant Keypoints. International Journal of Computer Vision 60, 2 (01 Nov 2004), 91–110. https: //doi.org/10.1023/B:VISI.0000029664.99615.94
- [14] B.S. Manjunath, C. Shekhar, and R. Chellappa. 1996. A new approach to image feature detection with applications. *Pattern Recognition* 29, 4 (1996), 627 – 640. https://doi.org/10.1016/0031-3203(95)00115-8
- [15] Yong-Jin Na, Eun-Soo; Jeong. 2013. FPGA Implementation of SURF-based Feature extraction and Descriptor generation. *Journal of Korea Multimedia Society* 16, 4 (2013), 483–492. https://doi.org/10.9717/kmms.2013.16.4.483
- [16] Adrien Bartoli Pablo Alcantarilla (Georgia Institute of Technolog), Jesus Nuevo (TrueVision Solutions AU). 2013. Fast Explicit Diffusion for Accelerated Features in Nonlinear Scale Spaces. In Proceedings of the British Machine Vision Conference. BMVA Press.
- [17] J. Rettkowski, A. Boutros, and D. Göhringer. 2015. Real-time pedestrian detection on a xilinx zynq using the HOG algorithm. In 2015 International Conference on ReConFigurable Computing and FPGAs (ReConFig). 1–8. https://doi.org/10.1109/ ReConFig.2015.7393339
- [18] E. Rosten and T. Drummond. 2005. Fusing points and lines for high performance tracking. In Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1, Vol. 2. 1508–1515 Vol. 2. https://doi.org/10.1109/ICCV.2005.104
- [19] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. 2011. ORB: An efficient alternative to SIFT or SURF. In 2011 International Conference on Computer Vision. 2564–2571. https://doi.org/10.1109/ICCV.2011.6126544