



- (51) International Patent Classification:
G06K 9/46 (2006.01) *G06K 9/62* (2006.01)
- (21) International Application Number:
PCT/IB2016/001890
- (22) International Filing Date:
15 December 2016 (15.12.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/973,272 17 December 2015 (17.12.2015) US
- (71) Applicant: **LINEAR ALGEBRA TECHNOLOGIES LIMITED** [IE/IE]; 19 Mountjoy Square East, Dublin 1 (IE).
- (72) Inventors: **MOLONEY, David**; 42 Iona Road, Glasnevin, Dublin 9 (IE). **DEGHANI, Alireza**; 83 The Links, Elm Park, Rock Road, Dublin 4 (IE).
- (81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM,

DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: SYSTEMS AND METHODS FOR PROVIDING AN IMAGE CLASSIFIER




FIG. 5

(57) Abstract: Systems and methods are provided for image classification using histograms of oriented gradients (HoG) in conjunction with a trainer. The efficiency of the process is greatly increased by first establishing a bitmap which identifies a subset of the pixels in the HoG window as including relevant foreground information, and limiting the HoG calculation and comparison process to only the pixels included in the bitmap.

SYSTEMS AND METHODS FOR PROVIDING AN IMAGE CLASSIFIER

Field of the application

[0001] This present application relates generally to computer vision processing, and more specifically to efficient classification of images.

Background

[0002] The field of computer vision often requires classifiers that are trained to detect objects such as faces and people, with a view to enabling applications that interact with people and real-world objects. A variety of classifiers exist, as computer vision researchers are consistently seeking more resource-efficient methods for accurately locating and identifying various objects in images.

[0003] One known method of identifying a particular class of object, described in FIG. 1, uses histograms of oriented gradients (HoG) in conjunction with training images and a learning system. HoG has been used to detect humans against a variety of backgrounds, as well as faces, animals, vehicles, and other objects. Because HoG uses a relatively compact reference descriptor, it has been successfully used in real-time to classify objects in streaming video. It has also been demonstrated to enable robust detection in the presence of rotations, scaling, and variations in terms of lighting conditions.

[0004] FIG. 1 illustrates a process 100 known in the art for classifying objects in images using HoG in conjunction with a support vector machine (SVM) algorithm – fittingly referred to in the art as HoG/SVM. The process as described uses the parameters identified by Dalal and Triggs in their 2005 paper: “Histograms of oriented gradients for human detection,” *International Conference on Computer Vision and Pattern Recognition*, Vol.2, pp. 886-893, June 2005, which is herein incorporated by reference in its entirety.

[0005] First, gradient values are calculated for each pixel within a particular cell (step 102 in FIG. 1). As shown in FIG. 2, which illustrates the process, a defined rectangular HoG detection window 202 is applied to a portion of the image, which divides the pixels into discrete cells 204. An HoG cell 204 may, for example, comprise 8 pixels on each side for a total of 64 (8-by-8) pixels per cell 204, although larger or smaller cell sizes may be chosen in some implementations. For each cell 204, a magnitude and orientation of a gradient is calculated. A variety of filters may be applied to calculate these values. For example, as one

implementation, the magnitude of the gradient $|G|$ may be given according to the intensity values of its adjacent pixels:

$$|G| = |G_x| + |G_y| = |G_{x+1} - G_{x-1}| + |G_{y+1} - G_{y-1}|.$$

And the orientation θ may be given according to the tangent of the horizontal x and vertical y intensities:

$$\theta = \arctan (|G_{y+1} - G_{y-1}| / |G_{x+1} - G_{x-1}|)$$

[0006] To create the histogram, the orientation angles θ are broken up into some number of bins. In this example, the range of 0° to 180° is broken into nine bins of 20° each. Each intensity value $|G|$ is added to the bin associated with its orientation angle θ (step 104 in FIG. 1). The resulting HoG cell descriptor, illustrated as 206 in FIG. 2, has 9 values each with a minimum of zero and a maximum of 128 times the maximum pixel intensity value.

[0007] Each of the cell descriptors 206 is then aggregated into block descriptors 210 (step 106 in FIG. 1), based on each 2-by-2 block 208 of four cells. Because every block 208 of cells is used, a cell 204 not on the edge of the window 202 will appear in four different blocks 208, and therefore its descriptor 206 will be included in four different block descriptors 212.

[0008] Each block descriptor 210, including the descriptors 206 of each of the four cells 204 in the block 208, is normalized according to the descriptors in that block (step 108 in FIG. 1). A variety of normalization algorithms can be used, many of which are discussed in the Dalal and Triggs 2005 paper referenced above. The result of this process is a normalized block descriptor 212 for each block 208, a set of histogram data representing 36 data elements per block. Because the normalization depends on the values of the four descriptors 206 in a particular block descriptor 210, the normalized values associated with a particular cell 206 may be different in each normalized block descriptor 212 that includes that cell.

[0009] For a 64-by-128 pixel window 202, the complete HoG descriptor 214 representing the normalized block descriptors 212 comprises 105 normalized blocks of histogram data: a total of 3,780 data values. This complete descriptor 214 is fed into the SVM classifier (step 110 in FIG. 1), which has previously evaluated training images according to the same parameters. The training images may be any appropriate set of training data for the objects being evaluated, such as the MIT and INRIA image data sets described in Dalal and Triggs 2005 paper. Other publicly available or proprietary training images can be used.

[0010] The HoG computation is performed by repeatedly stepping a window, 64 pixels wide by 128 pixels high across in the illustrated example, across a source image frame and

computing the HoG descriptor as outlined in the previous section. As the HoG calculation contains no intrinsic sense of scale and objects can occur at multiple scales within an image, the HoG calculation is stepped and repeated across each level of a scale pyramid.

[0011] FIG. 3 illustrates a window 302 being stepped across each level 304 of the scale pyramid 306. Each level 304 represents a further scaled-down copy of the image that is being scanned. The scaling factor between each level in the scale pyramid between one level and the next is commonly 1.05 or 1.2. The image is repeatedly down-scaled until the scaled source frame can no longer accommodate a complete HoG window.

[0012] The closed form for the number of pixels in the scaling pyramid is given by an expression based on s . s is the scale multiplier for the total number of pixels in the scale pyramid:

$$S = \frac{\alpha^{-m} - \alpha}{1 - \alpha}$$

α is the scaling factor used between pyramid levels. $m = \log(W/H)/\log(\alpha)$, where W and H are the respective width and height of the input image/video frame. The total number of pixels to consider in the scale pyramid is therefore $s*W*H$.

[0013] As can be seen in FIG. 2 and shown in the calculation above, the HoG descriptor for a system using 9D histograms for each of the 7-by-15 4-by-4 blocks in the image produces a 3.78kB descriptor for each 64-by-128 window that is examined in the upcoming image.

[0014] The images used to train such classifiers are typically rectangular as a by-product of the 2D image sensor arrays used to capture images. Add to this the simplicity to stepping a rectangular descriptor across a rectangular source image and convolving to detect a match and it is easy to see why this paradigm has taken root. While some objects, such as furniture, may indeed be square or rectangular, most objects of interest in classifying are not easily representable by simple geometric shapes. Therefore, a rectangular reference image is a poor match to such objects. Indeed, using a rectangular reference image means that significant additional work has to be done to convolve pixels that are not relevant to the matching task, and furthermore these pixels mean that some of the background surrounding the object of interest are aliased into the descriptor used to match images, thus confounding and degrading the accuracy of the matching operation.

[0015] The computational cost of each HoG data set is very high. One estimate is made by Dziri, Chevobbe, and Darouich in their 2013 paper: "Gesture recognition on smart

camera,” CEA LIST – Embedded Computing Laboratory, 2013. For example, to apply HoG to a 42-by-42 pixel region of interest requires the following operations: 11,664 addition, 1,296 multiplication, 5,200 division, 16 square root, and 5184 arctangent. The computation requires numerous costly and complex mathematical operations like division, square root, and arctangent, which take multiple cycles to implement on a conventional sequential processor in software. The computation also requires large numbers of more common mathematical operations like addition and multiplication, which typically execute in as little as one clock cycle. The computational costs are compounded by the fact that performing a brute-force search by stepping an HoG template for comparison over the entire image is even more computationally expensive depending on the resolution of the input image. Furthermore, in scenarios where objects may be seen at a range of distances, it is often necessary to search candidate windows of different sizes, further increasing the computational cost.

[0016] HoG/SVM is a very expensive operation. Many optimizations, from changing the scale factor to modifying the block-size in which the HoG window is stepped across the scaled source image, can be used to prune the search space and hence limit the computational effort. These factors combined mean that robust real-time HoG is confined to very high specification desktop systems that often offload computations to a high performance general processing unit (GPU). This pushes the power costs far beyond the bounds of mobile devices such as phones, tablets and mobile robots.

[0017] While it is possible to subsample the input image and perform a range of optimizations for mobile platforms, this often comes at a huge loss in terms of matching accuracy rendering the mobile implementation of very limited utility. Nonetheless, further optimizations to limit the computational expense of HoG processes are desired.

Summary

[0018] In accordance with the disclosed subject matter, systems and methods are provided for image classification using HoG in conjunction with a trainer. The efficiency of the process is greatly increased by first establishing a bitmap which identifies a subset of the pixels in the HoG window as including relevant foreground information, and limiting the HoG calculation and comparison process to only the pixels included in the bitmap.

[0019] The disclosed embodiments include a computer-implemented method for image classification. The method comprises: generating, from a set of training images, a non-

rectangular silhouette representing a relevant portion of an image window; generating, from the silhouette, a non-rectangular bitmap identifying fewer than all of the cells within a rectangular detection window, the rectangular detection window having a predetermined cell width and cell height; receiving an image for classification, the image including a cell width and cell height; applying the rectangular detection window to a portion of the image having the cell width and cell height of the rectangular detection window, wherein applying the detection window includes identifying cells within the image portion matching the identified cells of the non-rectangular bitmap; for each of the identified cells within the image portion, generating a histogram of orientations of the pixels within the cell; establishing overlapping cell blocks, each cell block having a plurality of contiguous identified cells from the image portion; generating, for each cell block, a descriptor including a histogram for each of the identified cells in the cell block, wherein each histogram within the descriptor is normalized to all of the histograms within the descriptor; and evaluating the generated descriptors against the set of training images in order to determine whether to identify the image portion as including a class of object associated with the set of training images.

[0020] In some embodiments, the method further comprises: re-scaling the image to a different cell width and cell height; and applying the rectangular detection window to a portion of the re-scaled image having the cell width and cell height of the rectangular detection window.

[0021] In some embodiments, the method further comprises generating the non-rectangular silhouette by one of averaging the set of training images, applying a computer vision algorithm, or applying deep learning to determine the relevant portion of the image window.

[0022] In some embodiments, the method further comprises generating the non-rectangular silhouette by: determining intensity values of all pixels in each position for each training image from the set of training images; adding the intensity values of all the pixels for each position from the set of training images to generate resulting values; and dividing the resulting values for each position from the set of training images by a number of training images from the set of training images,

[0023] In some embodiments, the method further comprises, for each of the identified cells within the image portion, generating a histogram of orientations of the pixels within the cell by: calculating gradients, each having a magnitude and a corresponding orientation;

associating each corresponding orientation to one of a plurality of bins; and adding each magnitude to the one of the plurality of bins associated with the corresponding orientation.

[0024] In some embodiments, the method further comprises providing each cell block as a two-by-two of the contiguous identified cells from the image portion.

[0025] In some embodiments, the method further comprises generating, for each cell block, a descriptor by concatenating the histogram for each of the identified cells in the cell block.

[0026] In some embodiments, the method further comprises determining that the class of object is one of: a person; a face; a non-rectangular geometric shape; or a non-geometric shape.

[0027] The disclosed embodiments include a computing device for image classification. The computing device comprises a memory configured to store a set of training images; a non-rectangular silhouette that represents a relevant portion of an image window from the set of training images; and a non-rectangular bitmap that identifies fewer than all of the cells within a rectangular detection window, the rectangular detection window having a predetermined cell width and cell height. The computing device also comprises a hardware accelerator coupled to the memory, the hardware accelerator configured to: receive as input an image for classification, the image including a cell width and cell height; apply the stored rectangular detection window to a portion of the image having the cell width and cell height of the rectangular detection window by identifying cells within the image portion that match the identified cells of the non-rectangular bitmap; for each of the identified cells within the image portion, generate and store a histogram of orientations of the pixels within the cell; and establish overlapping cell blocks, each cell block having a plurality of contiguous identified cells from the image portion. The computing device also comprises a processor coupled to the hardware accelerator and to the memory, the processor configured to: generate, for each cell block, a descriptor including a histogram for each of the identified cells in the cell block, wherein each histogram within the descriptor is normalized to all of the histograms within the descriptor; and evaluate the generated descriptors against the stored set of training images in order to determine whether to identify the image portion as including a class of object associated with the set of training images. The memory, hardware accelerator, and/or processor can be further configured to carry out any and all of the steps described with respect to any of the above embodiments.

[0028] The disclosed embodiments include a non-transitory computer readable medium. The non-transitory computer readable medium has executable instructions operable to cause a data processing apparatus to: generate, from a set of training images, a non-rectangular silhouette representing a relevant portion of an image window; generate, from the silhouette, a non-rectangular bitmap identifying fewer than all of the cells within a rectangular detection window, the rectangular detection window having a predetermined cell width and cell height; receive an image for classification, the image including a cell width and cell height; apply the rectangular detection window to a portion of the image having the cell width and cell height of the rectangular detection window by identifying cells within the image portion matching the identified cells of the non-rectangular bitmap; for each of the identified cells within the image portion, generate a histogram of orientations of the pixels within the cell; establish overlapping cell blocks, each cell block having a plurality of contiguous identified cells from the image portion; generate, for each cell block, a descriptor including a histogram for each of the identified cells in the cell block, wherein each histogram within the descriptor is normalized to all of the histograms within the descriptor; and evaluate the generated descriptors against the set of training images in order to determine whether to identify the image portion as including a class of object associated with the set of training images. The executable instructions can be further operable to cause the data processing apparatus to carry out any and all of the steps described with respect to any of the above embodiments.

[0029] The present invention will now be described in more detail with reference to particular embodiments thereof as shown in the accompanying drawings. While the present disclosure is described below with reference to particular embodiments, it should be understood that the present disclosure is not limited thereto. Those of ordinary skill in the art having access to the teachings herein will recognize additional implementations, modifications, and embodiments, as well as other fields of use, which are within the scope of the present disclosure as described herein, and with respect to which the present disclosure may be of significant utility.

Description of Drawings

[0030] Various objects, features, and advantages of the disclosed subject matter can be more fully appreciated with reference to the following detailed description of the disclosed subject matter when considered in connection with the following drawings, in which like reference numerals identify like elements. The accompanying figures are schematic and are

not intended to be drawn to scale. For purposes of clarity, not every component is labelled in every figure. Nor is every component of each embodiment of the disclosed subject matter shown where illustration is not necessary to allow those of ordinary skill in the art to understand the disclosed subject matter.

[0031] FIG. 1 is a flowchart describing an implementation of a HoG/SVM classification process as known in the art.

[0032] FIG. 2 illustrates an implementation of the HoG/SVM process as practiced absent a bitmap as known in the art.

[0033] FIG. 3 illustrates a pyramid scaling scheme as known in the art.

[0034] FIG. 4 illustrates a silhouette produced by averaging a set of training images in accordance with embodiments of the present invention.

[0035] FIG. 5 illustrates an implementation of the bHoG/SVM process applying a bitmap in accordance with embodiments of the present invention.

[0036] FIG. 6 is a flowchart describing an implementation of a HoG/SVM classification process in accordance with embodiments of the present invention.

[0037] FIG. 7 is a schematic depicting hardware configured in accordance with embodiments of the present invention.

[0038] FIG. 8 is a schematic depicting hardware configured in accordance with embodiments of the present invention.

[0039] FIG. 9 is a schematic depicting hardware configured in accordance with embodiments of the present invention.

[0040] FIG. 10 is a timing diagram depicting the operation bitmap scheduler for a typical bitmap in accordance with embodiments of the present invention.

[0041] FIG. 11 is a flowchart describing an implementation of a HoG/SVM classification process in accordance with embodiments of the present invention.

Detailed Description

[0042] In the following description, numerous specific details are set forth regarding the systems and methods of the disclosed subject matter and the environment in which such systems and methods may operate, etc., in order to provide a thorough understanding of the disclosed subject matter. It will be apparent to one skilled in the art, however, that the disclosed subject matter may be practiced without such specific details, and that certain features, which are well known in the art, are not described in detail in order to avoid

complication of the disclosed subject matter. In addition, it will be understood that the examples provided below are exemplary, and that it is contemplated that there are other systems and methods that are within the scope of the disclosed subject matter.

[0043] The present invention makes use of conditional processing using a bitmap which encodes those portions of an image used to train a classifier. The bitmap is derived by any suitable means. For example, the bitmap can be derived by averaging the images in a training data-set. As another example, the bitmap can be derived using computer vision algorithms or deep learning to determine a location and extent of an object to be detected in an image. These means produce a silhouette of the object, which is then thresholded in order to produce a bitmap. The bitmap may also take into account the underlying block structure used in the histogram of oriented gradients (HoG) computations.

[0044] FIG. 4 illustrates a plot having a silhouette generated from the average of the INRIA pedestrian detection image dataset, described with respect to the Dalal and Triggs 2005 paper referenced above. The plot shows a 48-by-96 pixel image 400, which is divided into 4-by-4 cells, for a total of 288 cells ($48/4\text{-by-}96/4 = 12\text{-by-}24 = 288$ 4-by-4 cells). Of the 288 cells, only 92 are included in the foreground silhouette 402 produced by averaging the images in the dataset, representing 32% ($92/288$) of the full area of the image.

[0045] The foreground silhouette can be generated in any number of suitable ways. For example, the silhouette can be generated by adding together the intensity values of all the pixels in each position within the training images, and then dividing the resulting values by the number of images. The following paper describes “average images” that can be generated by this process: Jun-Yan Zhu et al., “AverageExplorer: Interactive Exploration and Alignment of Visual Data Collections,” ACM Transactions on Graphics, Vol. 33, No. 4, August 2014. For example, the MATLAB functions `imadd()` and `imdivide()` can be used to perform these functions on a dataset and output the result to a bitmap file for review. The silhouette outline can be generated automatically based on the result of pixel-by-pixel averaging of the dataset, computer vision, deep learning, or other suitable means or combination of means.

[0046] The silhouette is then translated into an HoG detection bitmap 500 as shown in FIG. 5, in which the numbers 502-514 are similar to those used above, 202-214, with respect to the illustrated process of FIG. 2. Here, of the original 128 cells included in the HoG detection window 502, only 74 are included as foreground cells on the bitmap 500, representing 58% ($74/128$) of the total area of the window 502. Furthermore, in creating the

histogram data 506, six orientation bins of 30° each are used, instead of nine bins of 20° to further reduce the total size of the complete HoG descriptor 514. While covering the full detection window 502 requires 105 overlapping 2-by-2 blocks, covering the bitmap requires only 74 of such 2-by-2 blocks 508.

[0047] As shown in FIG. 5, the complete bHoG descriptor includes only 1776 total reference values (74-by-24-D), reducing the size of the descriptor by a factor of 3 compared to the full HoG descriptor described above. These descriptors are compared via SVM or another suitable learning algorithm to reference images that have been similarly reduced in size by the silhouette 400 of FIG. 4, yielding a further increase in efficiency.

[0048] By comparing only the relevant areas of the present image to only the relevant areas of the reference images, a significant reduction in process cost and complexity can be realized. In one experimental example, a bHoG/SVM classifier was trained using only those portions of the INRIA image database that did not contain background information and compared to a HoG/SVM classifier trained in the conventional manner. Results showed a 75% workload reduction using the INRIA image set for pedestrian detection, referenced above. A similar workload reduction was found using the Daimler Pedestrian Segmentation Benchmark Dataset introduced by F. Flohr and D. M. Gavrilu, "PedCut: an iterative framework for pedestrian segmentation combining shape models and multiple data cues," *Proc. of the British Machine Vision Conference*, Bristol, UK, 2013.

[0049] The results shown apply irrespective of the higher-level optimizations that can be made to reduce the dimensionality of the search space and are easily accelerated in hardware leading to a potential 4x speed-up in a hardware accelerated bHoG/SVM detector, with corresponding 4x reduction in memory storage, memory bandwidth and power consumption.

[0050] FIG. 6 is a flowchart describing an implementation of an HoG/SVM classification process as herein described. Process 600 starts at step 602 with an image to be classified. A detection window containing an arbitrarily-shaped bitmap of an arbitrarily-shaped silhouette is applied to a portion of the image. For example, a rectangular detection window containing a non-rectangular bitmap of a non-rectangular silhouette is applied to a portion of the image. The detection window is made up of cells and has a predetermined cell width and cell height. The bitmap identifies fewer than all of the cells within the regular detection window. The silhouette, and thus the bitmap, can be any suitable arbitrary shape, including any geometric shape (other than a rectangular shape), non-geometric shape, or combination of geometric and non-geometric shapes. The applying process identifies only the cells within the image

portion that match the identified cells of the bitmap. Next, at step 604, the process calculates gradient values (magnitudes and orientation angles) for each pixel within an identified cell within the image portion. At step 606, the process then breaks up the orientation angles into a suitable number of bins, and adds each magnitude to the bin associated with its orientation angle. Each resulting HoG cell descriptor is then aggregated into block descriptors at step 608. At step 610, each block descriptor, including its descriptors, are normalized using any suitable normalization algorithm. The resulting complete HoG descriptor, representing the normalized block descriptors, is then fed into the SVM classifier at step 110. In FIG. 6, steps 604, 606, 608, 610, and 612 each performs a similar function as described above in connection with respective steps 102, 104, 106, 108, and 110 in FIG. 1, but only on data corresponding to the identified cells of the bitmap corresponding to the silhouette, thereby resulting in optimizations, increased efficiencies, and a reduction in process cost and complexity as described above.

[0051] FIG. 7 shows a possible hardware block 700 for accelerated bHoG calculations as herein described. Input frames including an array of intensity values for a given cell and all adjacent cells are sent to a hardware bHoG accelerator 702, which includes modules for carrying out all necessary functions for computing gradients and generating orientation bins as earlier described. Certain modules may address and modify multiple HoG descriptors and associated bitmaps (that is, bHoG descriptors) stored in system memory 704. The remainder of the steps of the bHoG/SVM process (generating descriptor blocks, block normalization, and SVM classification of bHoG data) can be carried out by a VLIW processor 706 referencing the bHoG descriptors in system memory 704.

[0052] In FIG. 8, a possible hardware block to accelerate bHoG calculations and histogram binning is shown, where the angle calculations are performed and packed into memory for subsequent histogram binning in software. An integer divider 802 divides the θ angle computed for each pixel by the number of degrees associated with each histogram bin in order to output a bin number - an integer between 0 and N, where N is the number of histogram bins. This output is then used to select from the plurality of accumulators 804 in the histogram block to be updated on the next clock edge. The contents of these accumulators can then be read out at the end of the block of the bHoG calculation via a memory-mapped processor interface 806.

[0053] This hardware can be extended by the addition of a bitmap scheduler 900 to support sparse bitmaps as shown in FIG. 9. A bitmap counter 902 counts up to the max_count

limit – the number of bits in the bitmap corresponding to the selected cell-size. The output of the counter is the “count” and is used to look up a bit from the n-bit bitmap array 904 which is loaded via the bitmap input before circuit operation commences. The output of the bitmap array lookup 904 is the “HoG-enable” variable which is used to control the HoG computation block as well as the output of the “X_adr” and “Y_adr” outputs used to look up the required pixels for the HoG operation in memory by enabling the output AND gates 906 and 908. The inputs to these AND gates are generated by looking up “X_rel” and “Y_rel” relative addresses from a lookup-table 910 which is loaded before the circuit commences operation via the “lut_data” input. The “X_rel” and “Y_rel” addresses are added to the “X_base” and “Y_base” inputs via a pair of adders 912 and 914.

[0054] FIG. 10 is a timing diagram 1000 depicting the operation bitmap scheduler for an exemplary bitmap 1010. A bit stream 1012 is illustrated for the bitmap 1010, with the bits representing the first row of the bitmap 1010 marked on the bit stream 1012 as first row 1014. The portion of the timing diagram 1000 corresponding to the first row 1014 is marked as first timing portion 1002. As illustrated, a clock signal “clk” acts as a timer by sending alternating high and low signals as clock ticks. Within this context, various other signals representing different variables are sent. Most notably, and in accordance with the block diagram from FIG. 9, the “HoG_enable” signal as described above acts as a mask for the “X_adr” and “Y_adr” signals. If the “HoG_enable” signal is a high value, corresponding to a 1 value for the “bitmap” at that position, then the “X_adr” signal will be the sum of the “X_rel” and “X_base” signals and the “Y_adr” signal will be the sum of the “Y_rel” and “Y_base” signals. Else, if the “HoG_enable” signal is low, corresponding to a 0 value for the “bitmap” at that position, then the “X_adr” and “Y_adr” signals will also be low (that is, signaling a zero value or no value).

[0055] FIG. 11 is a flowchart describing an implementation of an HoG/SVM classification process as herein described. Process 1100 starts at step 1102 where the process generates, from a set of training images, an arbitrarily-shaped (such as a non-rectangular) silhouette representing a relevant portion of an image window. At step 1104, the process generates, from the silhouette, an arbitrarily-shaped (such as a non-rectangular) bitmap identifying fewer than all of the cells within a rectangular detection window, the rectangular detection window having a predetermined cell width and cell height. At step 1106, the process receives an image for classification, the image including a cell width and cell height. At step 1108, the process applies the rectangular detection window to a portion of the image

having the cell width and cell height of the rectangular detection window. This includes identifying cells within the image portion matching the identified cells of the non-rectangular bitmap. At step 1110, for each of the identified cells within the image portion, the process generates a histogram of orientations of the pixels within the cell. At step 1112, the process establishes overlapping cell blocks, each cell block having a plurality of contiguous identified cells from the image portion. At step 1114, the process generates, for each cell block, a descriptor including a histogram for each of the identified cells in the cell block. Each histogram within the descriptor is normalized to all of the histograms within the descriptor. At step 1116, the process evaluates the generated descriptors against the set of training images in order to determine whether to identify the image portion as including a class of object associated with the set of training images.

[0056] In order to minimize the memory requirements associated with bitmap storage, it could be compressed using run-length coding for each run of ones in each line of the bitmap if determined to be advantageous. In this case, a hardware and/or software run-length decoder would be used to reconstruct the bitmap from the run-length encoded version.

[0057] It will be appreciated that whilst several different arrangements have been described herein, that the features of each may be advantageously combined together in a variety of forms to achieve advantage.

[0058] In the foregoing specification, the application has been described with reference to specific examples. It will, however, be evident that various modifications and changes may be made therein without departing from the broader spirit and scope of the invention as set forth in the appended claims. For example, the connections may be any type of connection suitable to transfer signals from or to the respective nodes, units or devices, for example via intermediate devices. Accordingly, unless implied or stated otherwise the connections may for example be direct connections or indirect connections.

[0059] It is to be understood that the architectures depicted herein are merely exemplary, and that in fact many other architectures can be implemented which achieve the same functionality. In an abstract, but still definite sense, any arrangement of components to achieve the same functionality is effectively “associated” such that the desired functionality is achieved. Hence, any two components herein combined to achieve a particular functionality can be seen as “associated with” each other such that the desired functionality is achieved, irrespective of architectures or intermediate components. Likewise, any two components so

associated can also be viewed as being “operably connected,” or “operably coupled,” to each other to achieve the desired functionality.

[0060] Furthermore, those skilled in the art will recognize that boundaries between the functionality of the above described operations are merely illustrative. The functionality of multiple operations may be combined into a single operation, and/or the functionality of a single operation may be distributed in additional operations. Moreover, alternative embodiments may include multiple instances of a particular operation, and the order of operations may be altered in various other embodiments.

[0061] However, other modifications, variations and alternatives are also possible. The specifications and drawings are, accordingly, to be regarded in an illustrative rather than in a restrictive sense.

[0062] In the claims, any reference signs placed between parentheses shall not be construed as limiting the claim. The word “comprising” does not exclude the presence of other elements or steps than those listed in a claim. Furthermore, the terms “a” or “an,” as used herein, are defined as one or more than one. Also, the use of introductory phrases such as “at least one” and “one or more” in the claims should not be construed to imply that the introduction of another claim element by the indefinite articles “a” or “an” limits any particular claim containing such introduced claim element to inventions containing only one such element, even when the same claim includes the introductory phrases “one or more” or “at least one” and indefinite articles such as “a” or “an.” The same holds true for the use of definite articles. Unless stated otherwise, terms such as “first” and “second” are used to arbitrarily distinguish between the elements such terms describe. Thus, these terms are not necessarily intended to indicate temporal or other prioritization of such elements. The mere fact that certain measures are recited in mutually different claims does not indicate that a combination of these measures cannot be used to advantage.

CLAIMS

1. A computer-implemented method for image classification, comprising:
 - generating, from a set of training images, a non-rectangular silhouette representing a relevant portion of an image window;
 - generating, from the silhouette, a non-rectangular bitmap identifying fewer than all of the cells within a rectangular detection window, the rectangular detection window having a predetermined cell width and cell height;
 - receiving an image for classification, the image including a cell width and cell height;
 - applying the rectangular detection window to a portion of the image having the cell width and cell height of the rectangular detection window, wherein applying the detection window includes identifying cells within the image portion matching the identified cells of the non-rectangular bitmap;
 - for each of the identified cells within the image portion, generating a histogram of orientations of the pixels within the cell;
 - establishing overlapping cell blocks, each cell block having a plurality of contiguous identified cells from the image portion;
 - generating, for each cell block, a descriptor including a histogram for each of the identified cells in the cell block, wherein each histogram within the descriptor is normalized to all of the histograms within the descriptor; and
 - evaluating the generated descriptors against the set of training images in order to determine whether to identify the image portion as including a class of object associated with the set of training images.
2. The method of claim 1, further comprising:
 - re-scaling the image to a different cell width and cell height; and
 - applying the rectangular detection window to a portion of the re-scaled image having the cell width and cell height of the rectangular detection window.
3. The method of claim 1, wherein generating the non-rectangular silhouette comprises one or averaging the set of training images, applying a computer vision algorithm, or applying deep learning to determine the relevant portion of the image window.
4. The method of claim 1, wherein generating the non-rectangular silhouette comprises:
 - determining intensity values of all pixels in each position for each training image

from the set of training images;

adding the intensity values of all the pixels for each position from the set of training images to generate resulting values; and

dividing the resulting values for each position from the set of training images by a number of training images from the set of training images.

5. The method of claim 1, wherein, for each of the identified cells within the image portion, generating a histogram of orientations of the pixels within the cell, comprises:

calculating gradients, each having a magnitude and a corresponding orientation;

associating each corresponding orientation to one of a plurality of bins; and

adding each magnitude to the one of the plurality of bins associated with the corresponding orientation.

6. The method of claim 1, wherein each cell block comprises two-by-two of the contiguous identified cells from the image portion.

7. The method of claim 1, wherein generating, for each cell block, a descriptor comprises concatenating the histogram for each of the identified cells in the cell block.

8. The method of claim 1, further comprising determining that the class of object is one of:

a person;

a face;

a non-rectangular geometric shape; or

a non-geometric shape.

9. A computing device for image classification comprising:

a memory configured to store a set of training images, a non-rectangular silhouette that represents a relevant portion of an image window from the set of training images, and a non-rectangular bitmap that identifies fewer than all of the cells within a rectangular detection window, the rectangular detection window having a predetermined cell width and cell height;

a hardware accelerator coupled to the memory, the hardware accelerator configured to:

receive as input an image for classification, the image including a cell width and cell height,

apply the stored rectangular detection window to a portion of the image having the cell width and cell height of the rectangular detection window by identifying cells within the image portion that match the identified cells of the non-rectangular bitmap,

for each of the identified cells within the image portion, generate and store a histogram of orientations of the pixels within the cell, and

establish overlapping cell blocks, each cell block having a plurality of contiguous identified cells from the image portion; and

a processor coupled to the hardware accelerator and to the memory, the processor configured to:

generate, for each cell block, a descriptor including a histogram for each of the identified cells in the cell block, wherein each histogram within the descriptor is normalized to all of the histograms within the descriptor, and

evaluate the generated descriptors against the stored set of training images in order to determine whether to identify the image portion as including a class of object associated with the set of training images.

10. The computing device claim 9, wherein the hardware accelerator is further configured to:

re-scale the image to a different cell width and cell height; and

apply the rectangular detection window to a portion of the re-scaled image having the cell width and cell height of the rectangular detection window.

11. The computing device of claim 9, wherein at least one of the hardware accelerator or the processor is configured to:

average the set of training images; and

store the averaged set of training images as the non-rectangular silhouette.

12. The computing device of claim 9, wherein at least one of the hardware accelerator or the processor is configured to:

determine intensity values of all pixels in each position for each training image from the set of training images;

add the intensity values of all the pixels for each position from the set of training

images to generate resulting values;

divide the resulting values for each position from the set of training images by a number of training images from the set of training images; and

store the divided resulting values as the non-rectangular silhouette.

13. The computing device of claim 9, wherein the hardware accelerator configured to, for each of the identified cells within the image portion, generate and store a histogram of orientations of the pixels within the cell, is further configured to:

calculate gradients, each having a magnitude and a corresponding orientation;

associate each corresponding orientation to one of a plurality of bins; and

add each magnitude to the one of the plurality of bins associated with the corresponding orientation.

14. The computing device of claim 9, wherein the processor configured to evaluate the generated descriptors against the stored set of training images is further configured to determine whether the class of object is one of:

a person;

a face;

a non-rectangular geometric shape; or

a non-geometric shape.

15. A non-transitory computer readable medium having executable instructions operable to cause a data processing apparatus to:

generate, from a set of training images, a non-rectangular silhouette representing a relevant portion of an image window;

generate, from the silhouette, a non-rectangular bitmap identifying fewer than all of the cells within a rectangular detection window, the rectangular detection window having a predetermined cell width and cell height;

receive an image for classification, the image including a cell width and cell height;

apply the rectangular detection window to a portion of the image having the cell width and cell height of the rectangular detection window by identifying cells within the image portion matching the identified cells of the non-rectangular bitmap;

for each of the identified cells within the image portion, generate a histogram of orientations of the pixels within the cell;

establish overlapping cell blocks, each cell block having a plurality of contiguous identified cells from the image portion;

generate, for each cell block, a descriptor including a histogram for each of the identified cells in the cell block, wherein each histogram within the descriptor is normalized to all of the histograms within the descriptor; and

evaluate the generated descriptors against the set of training images in order to determine whether to identify the image portion as including a class of object associated with the set of training images.

16. The non-transitory computer readable medium of claim 15, further comprising executable instructions operable to cause the data processing apparatus to:

re-scale the image to a different cell width and cell height; and

apply the rectangular detection window to a portion of the re-scaled image having the cell width and cell height of the rectangular detection window.

17. The non-transitory computer readable medium of claim 15, further comprising executable instructions operable to cause the data processing apparatus to one of average the set of training images, apply a computer vision algorithm, or apply deep learning to generate the non-rectangular silhouette.

18. The non-transitory computer readable medium of claim 15, further comprising executable instructions operable to cause the data processing apparatus to:

determine intensity values of all pixels in each position for each training image from the set of training images;

add the intensity values of all the pixels for each position from the set of training images to generate resulting values; and

divide the resulting values for each position from the set of training images by a number of training images from the set of training images to generate the non-rectangular silhouette.

19. The non-transitory computer readable medium of claim 15, further comprising executable instructions operable to cause the data processing apparatus to:

calculate gradients, each having a magnitude and a corresponding orientation;

associate each corresponding orientation to one of a plurality of bins; and

add each magnitude to the one of the plurality of bins associated with the corresponding orientation.

20. The non-transitory computer readable medium of claim 15, further comprising executable instructions operable to cause the data processing apparatus to determine whether the class of object is one of:

- a person;
- a face;
- a non-rectangular geometric shape; or
- a non-geometric shape.

1/11


100


FIG. 1
PRIOR ART




Fig. 2
PRIOR ART




FIG. 3
PRIOR ART

4/11

400




FIG. 4




FIG. 5

6/11




FIG. 6




FIG. 7




FIG. 8

900





FIG. 9



11/11




FIG. 11

INTERNATIONAL SEARCH REPORT

International application No
PCT/IB2016/001890

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06K9/46 G06K9/62
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06K

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EP0-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	DALAL N ET AL: "Histograms of Oriented Gradients for Human Detection", PROCEEDINGS / 2005 IEEE COMPUTER SOCIETY CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, CVPR 2005 : [20 - 25 JUNE 2005, SAN DIEGO, CA], IEEE, PISCATAWAY, NJ, USA, vol. 1, 20 June 2005 (2005-06-20), pages 886-893, XP010817365, DOI: 10.1109/CVPR.2005.177 ISBN: 978-0-7695-2372-9 cited in the application SECTIONS 1, 4 AND 6 ----- -/--	1-20

☒ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

22 May 2017

Date of mailing of the international search report

31/05/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Miclea, Sorin

INTERNATIONAL SEARCH REPORT

International application No
PCT/IB2016/001890

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	JAIN SUYOG DUTT ET AL: "Predicting Sufficient Annotation Strength for Interactive Foreground Segmentation", 2013 IEEE INTERNATIONAL CONFERENCE ON COMPUTER VISION, IEEE, 1 December 2013 (2013-12-01), pages 1313-1320, XP032572784, ISSN: 1550-5499, DOI: 10.1109/ICCV.2013.166 [retrieved on 2014-02-28] Sections 2 and 3	1-20
A	MOHAMED ELMIKATY ET AL: "A novel two-level shape descriptor for pedestrian detection", SENSOR SIGNAL PROCESSING FOR DEFENCE (SSPD 2012), IEEE, 25 September 2012 (2012-09-25), pages 1-5, XP032433727, DOI: 10.1049/IC.2012.0122 ISBN: 978-1-84919-712-0 Section III	1-20
A	MOTTAGHI R: "Augmenting deformable part models with irregular-shaped object patches", COMPUTER VISION AND PATTERN RECOGNITION (CVPR), 2012 IEEE CONFERENCE ON, IEEE, 16 June 2012 (2012-06-16), pages 3116-3123, XP032232443, DOI: 10.1109/CVPR.2012.6248044 ISBN: 978-1-4673-1226-4 Sections 1 and 2	1-20
A	YA-LI HOU ET AL: "Multicue-Based Crowd Segmentation Using Appearance and Motion", IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS: SYSTEMS, IEEE, PISCATAWAY, NJ, USA, vol. 43, no. 2, 1 March 2013 (2013-03-01), pages 356-369, XP011493852, ISSN: 2168-2216, DOI: 10.1109/TSMCA.2012.2199308 Sections II and III	1-20
X,P	DEGHANI ALIREZA ET AL: "Object recognition speed improvement using BITMAP-HoG", 2016 IEEE INTERNATIONAL CONFERENCE ON IMAGE PROCESSING (ICIP), IEEE, 25 September 2016 (2016-09-25), pages 659-663, XP033016610, DOI: 10.1109/ICIP.2016.7532439 [retrieved on 2016-08-03] Section 3	1-20
	----- -/--	

INTERNATIONAL SEARCH REPORT

International application No
PCT/IB2016/001890

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2010/195899 A1 (NC PRAMOD [IN] ET AL) 5 August 2010 (2010-08-05) paragraphs [0068] - [0073] -----	1-20
A	MICHAEL PATZOLD ET AL: "Real-time person counting by propagating networks flows", ADVANCED VIDEO AND SIGNAL-BASED SURVEILLANCE (AVSS), 2011 8TH IEEE INTERNATIONAL CONFERENCE ON, IEEE, 30 August 2011 (2011-08-30), pages 66-70, XP032053725, DOI: 10.1109/AVSS.2011.6027296 ISBN: 978-1-4577-0844-2 Section 2 -----	1-20
A	LOPEZ-VILLA J S ET AL: "A novel tool for ground truth data generation for video-based object classification", 2015 20TH SYMPOSIUM ON SIGNAL PROCESSING, IMAGES AND COMPUTER VISION (STSIVA), IEEE, 2 September 2015 (2015-09-02), pages 1-6, XP032813500, DOI: 10.1109/STSIVA.2015.7330395 [retrieved on 2015-11-16] Sections 1 and 4 -----	1-20
A	ZHE LIN ET AL: "Shape-Based Human Detection and Segmentation via Hierarchical Part-Template Matching", IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, IEEE COMPUTER SOCIETY, USA, vol. 32, no. 4, 1 April 2010 (2010-04-01), pages 604-618, XP011287371, ISSN: 0162-8828, DOI: 10.1109/TPAMI.2009.204 Sections 2 and 3 -----	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/IB2016/001890

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010195899 A1	05-08-2010	CN 101799867 A US 2010195899 A1	11-08-2010 05-08-2010
