



(51) International Patent Classification:
G06K 9/32 (2006.01)

(21) International Application Number:
PCT/EP2017/060273

(22) International Filing Date:
28 April 2017 (28.04.2017)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicants: **TOYOTA MOTOR EUROPE** [BE/BE]; Avenue du Bourget 60, 1140 BRUSSELS (BE). **MAX-PLANCK-GESELLSCHAFT ZUR FÖRDERUNG DER WISSENSCHAFTEN E.V.** [DE/DE]; Hofgartenstraße 8, 80539 MUNCHEN (DE).

(72) Inventors: **REINO, Daniel Olmeda**; c/o TOYOTA MOTOR EUROPE, European Technical Administration Division, Avenue du Bourget 60, 1140 BRUSSELS (BE).

SCHIELE, Bernt; St. Ingberter Str. 39, 66125 SAARBRUCKEN (DE). **HOSANG, Jan Hendrik**; Großherzog-Friedrich Str. 116, 66121 SAARBRUCKEN (DE). **BENENSON, Rodrigo**; Riedweg 5, 8600 DUBENDORF (CH).

(74) Agent: **INTES, Didier** et al.; Cabinet BEAU DE LOMENIE, 158 rue de l'Université, 75340 PARIS CEDEX 07 (FR).

(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,

(54) Title: SYSTEM AND METHOD FOR DETECTING OBJECTS IN A DIGITAL IMAGE, AND SYSTEM AND METHOD FOR RESCORING OBJECT DETECTIONS.




FIG.3

(57) Abstract: The invention relates to a system for detecting objects in a digital image. The system comprises a neural network which is configured to generate candidate windows indicating object locations, and to generate for each candidate window a score representing the confidence of detection. Generating the scores comprises: - generating a latent representation for each candidate window, - updating the latent representation of each candidate window based on the latent representation of neighboring candidate windows, and - generating the score for each candidate window based on its updated latent representation. The invention further relates to a system for rescoring object detections in a digital image and to methods of detecting objects and rescoring objects.

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))

**System and method for detecting objects in a digital image, and
system and method for rescoring object detections**

FIELD OF THE DISCLOSURE

- 5 [1] The present disclosure is related to a system and a method for detecting objects in a digital image, and a system and a method for rescoring object detections.

BACKGROUND OF THE DISCLOSURE

- 10 [2] Modern object detectors follow a three step recipe: (1) proposing a search space of windows (exhaustive by sliding window or sparser using proposals), (2) scoring/ refining the window with a classifier/regressor, and (3) merging windows that might belong to the same object. This last stage is commonly referred to as "non-maximum suppression" (NMS), cf. e.g.:
- 15 [3] R. Girshick. Fast R-CNN. In ICCV, 2015,
 [4] P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. PAMI, 2010, and
 [5] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, and S. Reed. Ssd: Single shot multibox detector. In ECCV, 2016,
- 20 [6] L. Wan, D. Eigen, and R. Fergus. End-to-end integration of a convolutional network, deformable parts model and non-maximum suppression. In CVPR, 2015,
 [7] P. Henderson and V. Ferrari. End-to-end training of object class detectors for mean average precision. In ACCV, 2016.
- 25 [8] The de facto standard for such conventional NMS is a simple hand-crafted test time post-processing. The algorithm greedily selects high scoring detections and deletes close-by less confident neighbours since they are likely to cover the same object. This algorithm is simple, fast, and competitive compared to proposed alternatives.
- 30 [9] The most notable recent performance breakthrough in general object detection was marked by R-CNN, which effectively replaced features extraction and classifiers by a neural network, almost doubling performance on Pascal VOC, cf.:

[10] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In CVPR, 2014.

[11] Another significant improvement was to absorb the object proposal
5 generation into the network, cf. e.g.:

[12] S. Ren, K. He, R. Girshick, and J. Sun. Faster RCNN: Towards real-time object detection with region proposal networks. In NIPS, 2015.

[13] Meanwhile, other works avoid proposals altogether, leading to both speed and quality improvements, cf. e.g.:

10 [14] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In CVPR, 2016.

[15] There is a general trend towards end-to-end learning and it seems reasonable to expect further improvements by doing complete end-to-end training of detectors. NMS is one step in the pipeline that, for the most part,
15 has evaded the end-to-end learning paradigm. All of the above detectors train the classifier in a procedure that ignores the fact that the NMS problem exists and then runs a conventional NMS as a disconnected postprocessing.

[16] There is a need to overcome conventional NMS due to its significant conceptual shortcomings. Conventional NMS makes hard decision by deleting
20 detections and bases this decision on one fixed parameter that controls how wide the suppression is. A wide suppression would remove close-by high scoring detections that are likely to be false positives that hurt precision. On the other hand, if objects are close (e.g. in crowded scenes), close-by detections can be true positives, in which case suppression should be narrow to improve recall. When objects are close-by, conventional NMS is doomed to sacrifice
25 precision or recall independent of its parameter.

[17] There has further been an approach to combine decisions of a conventional NMS with different overlap thresholds, allowing the network to choose the conventional NMS operating point locally, cf:

30 [18] J Hosang, R. Benenson, and B. Schiele. A convent for non-maximum suppression. In GCPR, 2016.

SUMMARY OF THE DISCLOSURE

[19] Currently, it remains desirable to provide a system, rescoring system and method comprising a neural network which renders the need for conventional NMS post-processing superfluous.

5 [20] Therefore, according to the embodiments of the present disclosure, it is provided a system for detecting objects in a digital image. The system comprises a neural network being configured to generate candidate windows indicating object locations, and to generate for each candidate window a score representing the confidence of detection. Generating the scores (i.e. a score for
10 each candidate window) comprises:

- generating a latent representation for each candidate window,
- updating the latent representation of each candidate window based on the latent representation of neighboring candidate windows, and
- generating the score for each candidate window based on its updated
15 latent representation.

[21] Said neighboring candidate windows are referring desirably to the neighbors of the candidate window, whose latent representation is updated.

[22] Accordingly, the rescoring system is desirably configured to recognize multi detections by processing each object detection (i.e. generating
20 the latent representation of each candidate window) together with its neighboring detections (i.e. by updating each candidate window's representation, thereby considering the latent representations of the neighboring candidate windows of the currently updated candidate window). As a result, those object detections which are recognized to be multi detections of
25 the same object are desirably suppressed such that only one object detection remains (i.e. has a high score in relation to the its neighboring candidate windows).

[23] By providing such a system it becomes possible that the neural network learns NMS to overcome the limitations of a conventional NMS post-processing step. An NMS approach based on a neural network can learn to
30 adapt to the data distribution, overcome the tradeoff of conventional NMS, and importantly can be incorporated into a detector.

[24] A latent representation may be a multi-dimensional feature vector. For example, latent representation may be determined based on the candidate
35 window, and in particular image content in the candidate window. An example

would be to take the image content in the window and process it by at least one (or e.g. three) layer of the neural network, in order to obtain the latent representation/ feature vector. In a specific example the at least one layer consists of at least one of a convolution, a pooling, and/or a FC (fully
5 connected) layer.

[25] The latent representations may be updated such that the resulting scores of the candidate windows are changed.

[26] Desirably, this change is such that one candidate window of a detected object increases its score, while all other candidate windows on the
10 same object decrease their score. The amount of increase and/or decrease may be e.g. a learned amount that is deduced from the latent representations.

[27] Updating the latent representation of a candidate window may be performed by considering pairs of neighboring candidate windows. A pair of neighboring candidate windows may comprise said candidate window and one
15 of its neighboring candidate windows.

[28] The neighboring candidate windows may comprise all windows overlapping to a predetermined extent with the candidate window, of which the latent representation is updated.

[29] The neural network may comprise a repeating structure for updating
20 the latent representation of each candidate window based on the latent representation of neighboring candidate windows.

[30] Updating the latent representation of a candidate window may comprise forming with each of its neighboring candidate windows a pair of detections. The pair of detections may include the latent representations of said
25 candidate window and said neighboring window. The pair of detections may be a pair feature vector.

[31] Said pair feature vector may have equal or different dimensions as the feature vector of a latent representation. Generally, after every FC layer the dimensions may be freely selectable.

30 [32] Updating the latent representation of a candidate window may further comprise determining detection pair features based on the two candidate windows, for example, the geometry of the candidate windows.

[33] The detection pair features of the pair of candidate windows may be concatenated to the pair feature vector.

[34] The pair feature vector, to which the detection pair features are concatenated, may be mapped independently through fully connected layers.

[35] The number of pair feature vectors corresponding to the variable number of neighboring candidate windows may be reduced to a fixed-size representation by pooling. For example; the pair feature vectors may be reduced with an elementwise operation to one n-dimensional pooling feature vector.

[36] Said pooling feature vector may have equal or different dimensions as the feature vector of a latent representation and/or as the pair feature vector.

[37] Desirably the multiple (k) pair feature vectors of n dimensions are reduced to one n-dimensional pooling feature vector.

[38] The dimensionalities of the latent representations of candidate windows may be reduced before being combined into the pair feature vector.

[39] Accordingly, memory usage and/or computational cost can be reduced.

[40] Alternatively or additionally, after the pooling operation the dimensionalities may be increased to match the size of the latent representations of candidate windows.

[41] A candidate window may comprise a rectangular frame and/or a pixelwise mask of a detected object.

[42] The neural network may be trained by using at least one digital training image as a ground truth having a plurality of objects and respective object annotations indicating the actual locations of the objects. The training may comprise:

- generating candidate windows indicating object locations,
- for each candidate window generating a score representing the confidence of detection, and
- associating each object annotation to the best matching candidate window based on the score of the candidate windows and the overlap between the object annotation and the candidate windows, so that each candidate window is associated to maximally one object annotation and so that each object annotation is associated to maximally one candidate window.

[43] Accordingly, a matching (association) between unique couples of object annotations and candidate windows can be determined, so that none of the object annotations or the candidate window is matched (associated) twice.

[44] This matching may be obtained e.g. by:

- 5 - forming a list comprising the candidate windows, wherein the candidate windows are sorted in descending order by their confidence,
- identifying an unmatched object annotation having the highest overlap with the first one of the candidate windows in the list (i.e. having the highest confidence in the list),
- 10 - removing said candidate window from the list, and
- matching said candidate window to said object annotation, in case the overlap exceeds a predetermined threshold.

[45] The neural network may be trained by using successfully matched detections as positive training examples, and unmatched detections as negative
15 training examples.

[46] The invention further relates to a system for rescored object detections in a digital image. An object detection comprises a candidate window indicating the object location and a score representing the confidence of detection. The system comprises a neural network being configured to:

- 20 - generate a latent representation for each object detection,
- update the latent representation of each object detection based on the latent representation of neighboring object detections, and
- generate the new score for each object detection based on its updated latent representation.

[47] By providing such a system it becomes possible that the neural network learns NMS to overcome the limitations of a conventional NMS post-processing step. An NMS approach based on a neural network can learn to adapt to the data distribution, overcome the tradeoff of conventional NMS.

[48] Furthermore, by providing such a system the present disclosure
30 proposes a "pure NMS network" which is able to do the task of non-maximum suppression without image content or access to decisions of another algorithm.

[49] Said system for rescored object detections may be a part (or a sub-system) of the system for detecting objects in a digital image, as described above. In other words, the system for detecting objects may comprise the
35 system for rescored object detections.

[50] For example, the system for detecting objects in a digital image may comprise a neural network being configured to generate candidate windows indicating object locations, and for each candidate window generate a score representing the confidence of detection, wherein said neural system may
5 further comprise said system for rescored object detections, as described above.

[51] Alternatively, the system for detecting objects may comprise a first neural network being configured to generate candidate windows indicating object locations, and for each candidate window generate a score representing
10 the confidence of detection. Further, the system for detecting objects may comprise a second neural system constituting the system for rescored object detections.

[52] Desirably, this change is such that one candidate window of a detected object increases its score, while all other candidate windows on the
15 same object decrease their score. The amount of increase and/or decrease may be e.g. a learned amount that is deduced from the latent representations.

[53] Updating the latent representation of a candidate window may be performed by considering pairs of neighboring candidate windows. A pair of neighboring candidate windows may comprise said candidate window and one
20 of its neighboring candidate windows.

[54] The neighboring candidate windows may comprise all windows overlapping to a predetermined extent with the candidate window, of which the latent representation is updated.

[55] The neural network may comprise a repeating structure for updating
25 the latent representation of each candidate window based on the latent representation of neighboring candidate windows.

[56] Updating the latent representation of a candidate window may comprise forming with each of its neighboring candidate windows a pair of detections. The pair of detections may include the latent representations of said
30 candidate window and said neighboring window. The pair of detections may be a pair feature vector.

[57] Updating the latent representation of a candidate window may further comprise determining detection pair features based on the two candidate windows, for example, the geometry of the candidate windows.

[58] The detection pair features of the pair of candidate windows may be concatenated to the pair feature vector.

[59] The pair feature vector, to which the detection pair features are concatenated, may be mapped independently through fully connected layers.

5 [60] The number of pair feature vectors corresponding to the variable number of neighboring candidate windows may be reduced to a fixed-size representation by pooling. For example; the pair feature vectors may be reduced with an elementwise operation to one n-dimensional pooling feature vector.

10 [61] The dimensionalities of the latent representations of candidate windows may be reduced before being combined into the pair feature vector.

[62] Alternatively or additionally, after the pooling operation the dimensionalities may be increased to match the size of the latent representations of candidate windows.

15 [63] A candidate window may comprise a rectangular frame and/or a pixelwise mask of a detected object.

[64] The neural network may be trained by using at least one digital training image as a ground truth having a plurality of objects and respective object annotations indicating the actual locations of the objects. The training
20 may comprise:

- generating candidate windows indicating object locations,
- for each candidate window generating a score representing the confidence of detection, and
- associating each object annotation to the best matching candidate
25 window based on the score of the candidate windows and the overlap between the object annotation and the candidate windows, so that each candidate window is associating to maximally one object annotation.

[65] The neural network may be trained by using successfully matched detections as positive training examples, and unmatched detections as negative
30 training examples.

[66] The present disclosure further relates to a method of detecting objects in a digital image. In the method a neural network performs the steps of:

- generating candidate windows indicating object locations, and

- for each candidate window generating a score representing the confidence of detection,

The step of generating the scores comprises:

- generating a latent representation for each candidate window,
- 5 - updating the latent representation of each candidate window based on the latent representation of neighboring candidate windows, and
- generating the score for each candidate window based on its updated latent representation.

[67] The method may comprise further method steps which correspond to the functions of the system for detecting objects in a digital image, as described above. The further desirable method steps are described in the following.

[68] The present disclosure further relates to a method of rescore object detections in a digital image. An object detection comprises a candidate window indicating the object location and a score representing the confidence of detection. In the method a neural network performs the steps of:

- generating a latent representation for each object detection,
- updating the latent representation of each object detection based on the latent representation of neighboring object detections, and
- generating the new score for each object detection based on its updated latent representation.

[69] The method may comprise further method steps which correspond to the functions of the system for rescore object detections in a digital image, as described above. The further desirable method steps are described in the following.

[70] The present disclosure further relates to a computer program comprising instructions for executing the steps of the method of detecting objects in a digital image, when the program is executed by a computer.

[71] Finally, the present disclosure further relates to a computer program comprising instructions for executing the steps of the method of rescore object detections in a digital, when the program is executed by a computer.

[72] It is intended that combinations of the above-described elements and those within the specification may be made, except where otherwise contradictory.

[73] It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the disclosure, as claimed.

[74] The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate embodiments of the disclosure and together with the description, and serve to explain the principles thereof.

BRIEF DESCRIPTION OF THE DRAWINGS

[75] Fig. 1 shows a block diagram of a system with a detector and a rescoring system according to embodiments of the present disclosure;

[76] Fig. 2 shows a schematic diagram of how detection features are combined into a pairwise context according to embodiments of the present disclosure;

[77] Fig. 3 shows a schematic flow chart illustrating the operation of the rescoring system, in particular a neural network for the rescoring procedure according to embodiments of the present disclosure;

[78] Fig. 4a shows a schematic flow chart illustrating a training architecture of the neural network according to embodiments of the present disclosure; and

[79] Fig. 4b shows a schematic flow chart illustrating a test architecture of the neural network according to embodiments of the present disclosure.

DESCRIPTION OF THE EMBODIMENTS

[80] Reference will now be made in detail to exemplary embodiments of the disclosure, examples of which are illustrated in the accompanying drawings. Wherever possible, the same reference numbers will be used throughout the drawings to refer to the same or like parts.

[81] Fig. 1 shows a block diagram of a system 10 with an object detector 1 and a rescoring system 2 (i.e. a system for rescoring object detections) according to embodiments of the present disclosure. The system may have various further functions, e.g. may be a robotic system or a camera system. It may further be integrated in a vehicle.

[82] The system 10 may comprise an electronic circuit, a processor (shared, dedicated, or group), a combinational logic circuit, a memory that executes one or more software programs, and/or other suitable components

that provide the described functionality. In other words, system 10 may be a computer device. The system may be connected to a memory, which may store data, e.g. a computer program which when executed, carries out the method according to the present disclosure. In particular, the system or the memory
5 may store software which comprises the neural network according to the present disclosure.

[83] The system 10, in particular the detector 1, has an input for receiving a digital image or a stream of digital images. In particular, the system 10, in particular the detector 1, may be connected to an optical sensor 3, in
10 particular a digital camera. The digital camera 3 is configured such that it can record a scene, and in particular output digital data to the system 10, in particular the detector 1.

[84] The detector 1 may be implemented as software running on the system 10 or as a hardware element of the system 10. The detector 1 carries
15 out a computer vision algorithm for detecting the presence and location of objects in a sensed scene. For example, vehicles, persons, and other objects may be detected. The detector outputs candidate windows indicating object locations and generates for each candidate window a score representing the confidence of detection.

[85] Furthermore, also the rescoring system 2 may be implemented as a software running on the system 10 or as a hardware element of the system 10. In particular, the system may comprise a neural network which includes both the detector and the rescoring system. Alternatively the rescoring system may
20 be realized as an independent neural network (in particular beside a neural network comprising the detector).
25

[86] The rescoring system 2 receives as an input from the detector the detection results. In particular it receives information regarding one or more object detections. Each object detection comprises a candidate window indicating the object location and a score representing the confidence of
30 detection. The rescoring system rescoring these objects detections such that double detections are suppressed. In detail, the rescoring system generates a latent representation for each candidate window. The latent representation of each candidate window is subsequently updated based on the latent representation of neighboring candidate windows. The (thus re-evaluated)

score for each candidate window is then generated based on its updated latent representation.

[87] By this, the rescoring system is configured to recognize double detections by processing each object detection together with its neighboring
5 detections. Those object detections which are recognized to be multi detections of the same object are suppressed such that only one object detection remains.

[88] In the following the operation of the neural network according to the present disclosure is explained in more detail and with reference to figures 2 to 4b.

10 [89] Present-day detectors do not return all detections that have been scored, but instead use NMS as a post-processing step to remove redundant detections. In order to have true end-to-end learned detectors, the present disclosure relates to detectors without any post-processing. To understand why NMS is necessary, it is useful to look at the task of detection and how it is
15 evaluated.

[90] The task of object detection is to map an image to a set of boxes (i.e. candidate windows): one box per object of interest in the image, each box tightly enclosing an object. This means detectors ought to return exactly one detection per object. Since uncertainty is an inherent part of the detection
20 process, evaluations allow detections to be associated to a confidence. Confident erroneous detections are penalized more than less confident ones. In particular mistakes that are less confident than the least confident correct detection are not penalized at all.

[91] The detection problem can be interpreted as a classification problem
25 that estimates probabilities of object classes being present for every possible detection in an image. This viewpoint gives rise to "hypothesize and score" detectors that build a search space of detections (e.g. sliding window, proposals) and estimate class probabilities independently for each detection. As a result, two strongly overlapping windows covering the same object will both
30 result in high score since they look at almost identical image content. In general, instead of one detection per object, each object triggers several detections of varying confidence, depending on how well the detection windows cover the object.

[92] Since the actual goal is to generate exactly one detection per object
35 (or exactly one high confidence detection), a common practice (since at least

1994) is to assume that highly overlapping detections belong to the same object and collapse them into one detection. The predominant algorithm (conventional NMS) accepts the highest scoring detection, then rejects all detections that overlap more than some threshold and repeats the procedure
5 with the remaining detections, i.e. greedily accepting local maxima and discarding their neighbours, hence the name. This algorithm eventually also accepts wrong detections, which is no problem if their confidence is lower than the confidence of correct detections.

[93] This conventional NMS algorithm works well if (1) the suppression is
10 wide enough to always suppress high scoring detections triggered by same object and (2) the suppression is narrow enough to never suppress high scoring detections of the next closest object. If objects are far apart condition (2) is easy to satisfy and a wide suppression works well. In crowded scenes with high occlusion between objects there is a tension between wide and narrow
15 suppression. In other words with one object per image NMS is trivial, but highly occluded objects require a better NMS algorithm.

[94] The present disclosure is based on these requirements to an NMS algorithm but proposes a solution where the NMS is a "pure NMS network", in particular it can be incorporated into a detector.

20 [95] Independent processing of image windows leads to overlapping detection giving similar scores, this is a requirement of robust functions: similar inputs lead to similar outputs. A detector that outputs only one high scoring detection per object thus has to be also conditioned on other detections: multiple detections on the same object should be processed jointly, so the
25 detector can tell there are repeated detections and only one of them should receive a high score.

[96] Typical inference of detectors consist of a classifier that discriminates between image content that contains an object and image content that does not. The positive and negative training examples for this detector are usually
30 defined by some measure of overlap between objects and bounding boxes. Since similar boxes will produce similar confidences anyway, small perturbation of object locations can be considered positive examples, too. This technique augments the training data and leads to more robust detectors. Using this type of classifier training does not reward one high scoring detection per object, and
35 instead deliberately encourages multiple high scoring detections per object.

[97] From this analysis there result two key ingredients for the system of the present disclosure, in order for a detector to generate exactly one detection per object:

5 [98] 1. A loss that penalises double detections to teach the detector, as precisely one detection per object is required.

[99] 2. Joint processing of neighbouring detections so the detector has the necessary information to tell whether an object was detected multiple times.

10 [100] The neural network design according to the present disclosure accommodates both ingredients. The neural network design avoids hard decisions and does not discard detections to produce a smaller set of detections. Instead, NMS is reformulated as a rescoring task that seeks to decrease the score of detections that cover objects that already have been detected. After rescoring, simple thresholding may be done to reduce the set of
15 detections. For evaluation the full set of rescored detections may be passed to the evaluation script without any post processing.

Loss:

20 [101] A detector is supposed to output exactly one high scoring detection per object. The loss for such a detector must inhibit multiple detections of the same object, irrespective of how close these detections are.

[102] The detector may be judged by the evaluation criterion of a benchmark, which in turn defines a matching strategy to decide which
25 detections are correct or wrong. This is the matching that should be used at training time. Typically benchmarks sort detections in descending order by their confidence and match detections in this order to objects, preferring most overlapping objects. Since already matched objects cannot be matched again surplus detections are counted as false positives that decrease the precision of
30 the detector. This matching strategy may be used for training.

[103] The result of the matching may be used as labels for the classifier: successfully matched detections are positive training examples, while unmatched detections are negative training examples for a standard binary loss. Typically all detections that are used for training of a classifier have a label
35 associated as they are fed into the network. In this case the network has

access to detections and object annotations and the matching layer generates labels, that depend on the predictions of the network. This class assignment directly encourages the rescoreing behaviour that is desired to be achieved.

[104] Let d_i denote a detection, y_i (as element of $[-1, 1]$) indicate whether
 5 or not d_i was successfully matched to an object, and let f denote the scoring function that jointly scores all detections on an image

$$f([d_i]_{i=1}^n) = [s_i]_{i=1}^n.$$

[105] It is trained with weighted loss
 10

$$L(s_i, y_i) = \sum_{i=1}^N w_{y_i} \cdot \log(1 + \exp(-s_i \cdot y_i)).$$

[106] Here loss per detection is coupled to the other detections through the matching that produces y_i . The weighting w_{y_i} is used to counteract the extreme class imbalance of the detection task. The weights may be chosen so
 15 the expected class conditional weight of an example equals a parameter

$$\mathbb{E}(w_1 I(y_i = 1)) = \gamma.$$

[107] When generalising to the multiclass setting, detections are associated to both a confidence and a class. Since only detections are rescored, detections are allowed to be "switched off" but not to change their class. As a result, only detections are matched to objects of the same class, but the
 20 classification problem remains binary and the above loss still applies. When representing the detection scores, a one-hot encoding may be used: a zero vector that only contains the score at the location in the vector that corresponds to the class. Since mAP computation does not weight classes by their size, the instance weights may be assigned in a way that their expected
 25 class conditional weight is uniformly distributed.

Joint processing:

[108] In order to effectively minimize the aforementioned loss, it is necessary for the neural network to jointly process detections. To this end a
5 neural network is designed with a repeating structure, which is called "blocks" (shown in figure 3). One block gives each detection access to the representation of its neighbours and subsequently updates its own representation. Stacking multiple blocks means the network alternates between allowing every detection "talk" to its neighbours and updating its own
10 representation. In other words, detections talk to their neighbours to update their representation.

[109] There are two non-standard operations here that are key. The first is a layer, that builds representations for pairs of detections, as shown in figure 2. This leads to the key problem: an irregular number of neighbours for each
15 detection. Since it is desired to avoid a discretisation scheme, this issue is desirably solved with pooling across detections (the second key).

Detection features:

[110] The blocks of the neural network take the detection feature vector of
20 each detection as input and outputs an updated vector (see high-level illustration in figure 4a, 4b). Outputs from one block are input to the next one. The values inside this $c = 128$ dimensional feature vector are learned implicitly during the training. The output of the last block is used to generate the new
25 detection score for each detection.

[111] The first block takes an all-zero vector as input. The detections' information is fed into the network in the "pairwise computations" section of figure 3 as described below. This zero input could potentially be replaced with
30 image features.

Pairwise detection context:

[112] The first is a layer that builds representations for pairs of detections, as shown in figure 2. Fig. 2 shows a schematic diagram of how detection
35 features are combined into a pairwise context according to embodiments of the

present disclosure. Each solid block is the feature vector of the detection of corresponding pattern (e.g. different dashed lines). The hatched blocks are the "detection pair features" that are defined by the two detections corresponding to the two patterns.

- 5 [113] Each mini-batch consists of all n detections on an image, each represented by a c dimensional feature vector, so the data has size $n \times c$ and accessing to another detection's representations means operating within the batch elements. A detection context layer is used, that, for every detection d_i , generates all pairs of detections $(d_i; d_j)$ for which d_j sufficiently overlaps with d_i (IoU > 0.2). The representation of a pair of detections consists of the
 10 concatenation of both detection representations and g dimensional detection pair features (see below), which yields an $l = 2c + g$ dimensional feature. To process each pair of detections independently, the features are arranged of all pairs of detections along the batch dimension: if detection d_i has k_i
 15 neighbouring detection that yields a batch of size $K \times l$, where

$$K = \sum_{i=1}^n (k_i + 1)$$

- since also the pair $(d_i; d_i)$ is included. The number of neighbours k_i (the number of pairs) is different for every detection even within one mini-batch. To reduce the variable sized neighbourhood into a fixed size representation, the
 20 architecture of the present disclosure uses global max-pooling over all detection pairs that belong to the same detection ($K \times l \rightarrow n \times l$), after which normal fully connected layers can be used to update the detection representation (see figure 3).

- 25 Detection pair features:

- [114] The features for each detection pair used in the detection context consists of several properties of a detection pair: (1) the intersection over union (IoU), (2-4) the normalised distance in x and y direction and the normalised L_2
 30 distance (normalized by the average of width and height of the detection), (4-5) scale difference of width and height (e.g. $\log(w_i/w_j)$), (6) aspect ratio difference $\log(a_i/a_j)$, (7-8) the detection scores of both detections. In the multi-class setup, each detection provides a scores vector instead of a scalar thus increasing the number of pair features. All these raw features are fed into

3 fully connected layers, to learn the g detection pair features that are used in each block.

Block:

5

[115] A block does one iteration allowing detections to look at their respective neighbours and updating their representation, as shown in figure 3. Fig. 3 shows a schematic flow chart illustrating the operation of the rescoring system, in particular a neural network for the updating procedure according to
10 embodiments of the present disclosure. One block of the neural network of the present disclosure is shown here for one detection. The representation of each detection is reduced and then combined into neighbouring detection pairs and concatenated with detection pair features (hatched boxes, corresponding features and detections have the same pattern). Features of detection pairs are
15 mapped independently through fully connected layers. The variable number of pairs is reduced to a fixed-size representation by max-pooling. Pairwise computations are done for each detection independently.

[116] In particular, the neural network consists of a dimensionality reduction, a pairwise detection context layer, 2 fully connected layers applied to
20 each pair independently, pooling across detections, and two fully connected layers, where the last one increases dimensionality again. The input and output of a block are added as in the Resnet architecture, cf.:

[117] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In ECCV, 2016

25 [118] The first block receives zero features as inputs, so all information that is used to make the decision is bootstrapped from the detection pair features. The output of the last block is used by three fully connected layers to predict a new score for each detection independently (cf. figure 4a, 4b).

[119] Fig. 4a shows a schematic flow chart illustrating a training
30 architecture of the neural network according to embodiments of the present disclosure; and Fig. 4b shows a schematic flow chart illustrating a test architecture of the neural network according to embodiments of the present disclosure. Fig. 4a and 4b are understood to be high level diagram of the disclosure. In the training architecture of fig. 4a blocks are, as described in
35 figure 3. "FC" denotes fully connected layers. All features in this diagram may

have 128 dimensions (input vector and features between the layers/blocks). The output is a scalar.

Parameters:

5

[120] The neural networks may have 16 blocks. The feature dimension for the detection features may be 128 and may be reduced to 32 before building the pairwise detection context. The detection pair features may also have 32 dimensions. The fully connected layers after the last block may output 128
10 dimensional features. When the feature dimension is changed, the ratio between the number of features in each layer is kept constant, so indicating the detection feature dimension is sufficient.

Message passing:

15

[121] The forward pass over serveral stacked blocks can be interpreted as message passing. Every detection sends messages to all of its neighbours in order to negotiate which detection is assigned an object and which detections should decrease their scores. Instead of handcrafting the message passing
20 algorithm and its rules, the network is configured to latently learn the messages that are being passed.

[122] Throughout the description, including the claims, the term "comprising a" should be understood as being synonymous with "comprising at least one" unless otherwise stated. In addition, any range set forth in the
25 description, including the claims should be understood as including its end value(s) unless otherwise stated. Specific values for described elements should be understood to be within accepted manufacturing or industry tolerances known to one of skill in the art, and any use of the terms "substantially" and/or "approximately" and/or "generally" should be understood to mean falling within
30 such accepted tolerances.

[123] Although the present disclosure herein has been described with reference to particular embodiments, it is to be understood that these embodiments are merely illustrative of the principles and applications of the present disclosure.

[124] It is intended that the specification and examples be considered as exemplary only, with a true scope of the disclosure being indicated by the following claims.

CLAIMS

1. A system for detecting objects in a digital image, the system comprising a neural network being configured to:
 - 5 generate candidate windows indicating object locations, and for each candidate window generate a score representing the confidence of detection,wherein generating the scores comprises:
 - generating a latent representation for each candidate window,
 - 10 - updating the latent representation of each candidate window based on the latent representation of neighboring candidate windows, and
 - generating the score for each candidate window based on its updated latent representation.
- 15 2. The system according to claim 1, wherein the latent representations are updated such that the resulting scores of the candidate windows are changed.
3. The system according to any one of the preceding claims, wherein
 - 20 updating the latent representation of a candidate window is performed by considering pairs of neighboring candidate windows, a pair of neighboring candidate windows comprising said candidate window and one of its neighboring candidate windows.
- 25 4. The system according to any one of the preceding claims, wherein the neighboring candidate windows comprise all windows overlapping to a predetermined extent with the candidate window, of which the latent representation is updated.
- 30 5. The system according to any one of the preceding claims, wherein the neural network comprises a repeating structure for updating the latent representation of each candidate window based on the latent representation of neighboring candidate windows.

6. The system according to any one of the preceding claims, wherein updating the latent representation of a candidate window comprises forming with each of its neighboring candidate windows a pair of detections, the pair of detections including the latent representations of said candidate
5 window and said neighboring window,
the pair of detections being a pair feature vector.

7. The system according to the preceding claim, wherein updating the latent representation of a candidate window further comprises determining
10 detection pair features based on the two candidate windows, in particular the geometry of the candidate windows,
the detection pair features of the pair of candidate windows are concatenated to the pair feature vector.

15 8. The system according to any one of the preceding claims 6 or 7, wherein the pair feature vector, to which the detection pair features are concatenated, is mapped independently through fully connected layers.

9. The system according to any one of the preceding claims 6 to 8,
20 wherein the number of pair feature vectors corresponding to the variable number of neighboring candidate windows is reduced to a fixed-size representation by pooling, wherein the pair feature vectors are reduced with an elementwise operation to one n-dimensional pooling feature vector.

25 10. The system according to the preceding claim, wherein the dimensionalities of the latent representations of candidate windows are reduced before being combined into the pair feature vector, and/or
after the pooling operation the dimensionalities are increased to match the size of the latent representations of candidate windows.

30

11. The system according to any one of the preceding claims, wherein a candidate window comprises a rectangular frame and/or a pixelwise mask of a detected object.

12. The system according to any one of the preceding claims, wherein the neural network is trained by using at least one digital training image as a ground truth having a plurality of objects and respective object annotations indicating the actual locations of the objects, the training comprising:
- 5 generating candidate windows indicating object locations,
for each candidate window generating a score representing the confidence of detection, and
associating each object annotation to the best matching candidate window based on the score of the candidate windows and the overlap between the
10 object annotation and the candidate windows, so that each candidate window is associated to maximally one object annotation and so that each object annotation is associated to maximally one candidate window.
13. The system according to the preceding claim, wherein the neural
15 network is trained by using successfully matched detections as positive training examples, and unmatched detections as negative training examples.
14. A system for rescored object detections in a digital image, wherein an object detection comprises a candidate window indicating the object location
20 and a score representing the confidence of detection,
the system comprising a neural network being configured to:
- generate a latent representation for each object detection,
 - update the latent representation of each object detection based on the latent representation of neighboring object detections, and
 - 25 - generate the new score for each object detection based on its updated latent representation.
15. A method of detecting objects in a digital image,
a neural network performs the steps of:
- 30 generating candidate windows indicating object locations, and
for each candidate window generating a score representing the confidence of detection,
wherein the step of generating the scores comprises:
- generating a latent representation for each candidate window,

- updating the latent representation of each candidate window based on the latent representation of neighboring candidate windows, and generating the score for each candidate window based on its updated latent representation.

5

16. A method of rescoring object detections in a digital image, wherein an object detection comprises a candidate window indicating the object location and a score representing the confidence of detection, wherein a neural network performs the steps of:



10

- generating a latent representation for each object detection,
- updating the latent representation of each object detection based on the latent representation of neighboring object detections, and
- generating the new score for each object detection based on its updated latent representation.

15

17. A computer program comprising instructions for executing the steps of the method according to claim 15 or claim 16 when the program is executed by a computer.

1/3









FIG.2





INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2017/060273

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06K9/32
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)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Jan Hosang: "GitHub Hosang/Gossipnet", 13 February 2017 (2017-02-13), XP055435463, Retrieved from the Internet: URL:https://github.com/hosang/gossipnet/tree/ce8ea8d4aa7fadcd749a0ebd290b1574d348cd9d [retrieved on 2017-12-15] the whole document ----- -/--	1-17



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

15 December 2017

Date of mailing of the international search report

02/01/2018

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

Alecu, Teodor Iulian

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2017/060273

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
L	Jan Hosang ET AL: "Learning non-maximum suppression", 9 May 2017 (2017-05-09), XP55435200, Retrieved from the Internet: URL:https://arxiv.org/pdf/1705.02950.pdf [retrieved on 2017-12-14] the whole document	
X	HOSANG JAN ET AL: "A Convnet for Non-maximum Suppression", 27 August 2016 (2016-08-27), NETWORK AND PARALLEL COMPUTING; [LECTURE NOTES IN COMPUTER SCIENCE; LECT.NOTES COMPUTER], SPRINGER INTERNATIONAL PUBLISHING, CHAM, PAGE(S) 192 - 204, XP047353657, ISSN: 0302-9743 ISBN: 978-3-540-28012-5 [retrieved on 2016-08-27] Sections 1 and 2	1,2,4,5, 11-17
A	MATTHEW B BLASCHKO ET AL: "Non Maximal Suppression in Cascaded Ranking Models", 17 June 2013 (2013-06-17), IMAGE ANALYSIS, SPRINGER BERLIN HEIDELBERG, BERLIN, HEIDELBERG, PAGE(S) 408 - 419, XP047030599, ISBN: 978-3-642-38885-9 Section 3	1-17
A	RUJIKIETGUMJORN SITAPA ET AL: "Optimized Pedestrian Detection for Multiple and Occluded People", IEEE COMPUTER SOCIETY CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION. PROCEEDINGS, IEEE COMPUTER SOCIETY, US, 23 June 2013 (2013-06-23), pages 3690-3697, XP032493190, ISSN: 1063-6919, DOI: 10.1109/CVPR.2013.473 [retrieved on 2013-10-02] Sections 1-3	1-17
A	Rasmus Rothe ET AL: "Non-maximum suppression for object detection by passing messages between windows", 1 January 2014 (2014-01-01), XP55435204, DOI: 10.3929/ethz-a-010254996 Retrieved from the Internet: URL:http://e-collection.library.ethz.ch/eserv/eth:46736/eth-46736-01.pdf Sections 1- 3	1-17