



(51) International Patent Classification:

G06F 3/01 (2006.01) G06N 3/04 (2006.01)
G06F 3/0346 (2013.01) G06T 7/215 (2017.01)
G06K 9/66 (2006.01)

(21) International Application Number:

PCT/US2017/061618

(22) International Filing Date:

14 November 2017 (14.11.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/422,547 15 November 2016 (15.11.2016) US

(71) Applicant: **MAGIC LEAP, INC.** [US/US]; 7500 W. Sunrise Blvd., Plantation, FL 33322 (US).

(72) Inventors: **MALISIEWICZ, Tomasz**; 7500 W. Sunrise Blvd., Plantation, FL 33322 (US). **RABINOVICH, Andrew**; 7500 W. Sunrise Blvd., Plantation, FL 33322 (US). **BADRINARAYANAN, Vijay**; 7500 W. Sunrise Blvd., Plantation, FL 33322 (US). **DWIBEDI, Debidatta**; 1400 Wightman Street, Pittsburgh, PA 15217 (US).

(74) Agent: **JUANG, Agnes**; Knobbe, Martens, Olson & Bear LLP, 2040 Main Street, 14th Floor, Irvine, CA 92614 (US).

(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, JO, 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,

(54) Title: DEEP LEARNING SYSTEM FOR CUBOID DETECTION




FIG. 2

(57) Abstract: Systems and methods for cuboid detection and keypoint localization in images are disclosed. In one aspect, a deep cuboid detector can be used for simultaneous cuboid detection and keypoint localization in monocular images. The deep cuboid detector can include a plurality of convolutional layers and non-convolutional layers of a trained convolution neural network for determining a convolutional feature map from an input image. A region proposal network of the deep cuboid detector can determine a bounding box surrounding a cuboid in the image using the convolutional feature map. The pooling layer and regressor layers of the deep cuboid detector can implement iterative feature pooling for determining a refined bounding box and a parameterized representation of the cuboid.



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))

DEEP LEARNING SYSTEM FOR CUBOID DETECTION

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of priority to U.S. Patent Application Number 62/422,547, filed November 15, 2016, entitled “DEEP CUBOID DETECTION: BEYOND 2D BOUNDING BOXES,” the content of which is hereby incorporated by reference herein in its entirety.

BACKGROUND

Field

[0002] The present disclosure relates generally to systems and methods for three-dimensional object detection in images and more particularly to deep machine learning systems for detecting cuboids in images.

Description of the Related Art

[0003] A deep neural network (DNN) is a computation machine learning method. DNNs belong to a class of artificial neural networks (NN). With NNs, a computational graph is constructed which imitates the features of a biological neural network. The biological neural network includes features salient for computation and responsible for many of the capabilities of a biological system that may otherwise be difficult to capture through other methods. In some implementations, such networks are arranged into a sequential layered structure in which connections are unidirectional. For example, outputs of artificial neurons of a particular layer can be connected to inputs of artificial neurons of a subsequent layer. A DNN can be a NN with a large number of layers (e.g., 10s, 100s, or more layers).

[0004] Different NNs are different from one another in different perspectives. For example, the topologies or architectures (e.g., the number of layers and how the layers are interconnected) and the weights of different NNs can be different. A weight can be approximately analogous to the synaptic strength of a neural connection in a biological system. Weights affect the strength of effect propagated from one layer to another. The

output of an artificial neuron can be a nonlinear function of the weighted sum of its inputs. The weights of a NN can be the weights that appear in these summations.

SUMMARY

[0005] Building a three-dimensional (3D) representation of the world from a single monocular image is an important challenge in computer vision. The present disclosure provides examples of systems and methods for detection of 3D cuboids (e.g., box-like objects) and localization of keypoints in images. In one aspect, a deep cuboid detector can be used for simultaneous cuboid detection and keypoint localization in images. The deep cuboid detector can include a plurality of convolutional layers and non-convolutional layers of a trained convolutional neural network for determining a convolutional feature map from an input image. A region proposal network of the deep cuboid detector can determine a bounding box surrounding a cuboid in the image using the convolutional feature map. The pooling layer and regressor layers of the deep cuboid detector can implement iterative feature pooling for determining a refined bounding box and a parameterized representation of the cuboid.

[0006] Details of one or more implementations of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages will become apparent from the description, the drawings, and the claims. Neither this summary nor the following detailed description purports to define or limit the scope of the inventive subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1A is an example monocular image illustrating two-dimensional (2D) object detection with a bounding box overlaid around an object detected.

[0008] FIG. 1B is an example monocular image illustrating three-dimensional (3D) cuboid detection with a representation of the cuboid overlaid on the object detected. FIG. 1B shows that one cuboid inside the monocular image is detected and its vertices localized (shown as eight black circles that are connected).

[0009] FIG. 2 depicts an example architecture of a cuboid detector.

[0010] FIG. 3 is an example image illustrating region of interest (RoI) normalized coordinates.

[0011] FIGS. 4A-4G show images illustrating example cuboid detection and keypoint localization. One or more cuboids have been detected in each image with keypoint of each cuboid localized, shown as white connected circles.

[0012] FIGS. 5A-5C show example images showing improved performance with keypoint refinement via iterative feature pooling.

[0013] FIG. 6 is a schematic illustration show example cuboid vanishing points.

[0014] FIGS. 7A-7F are plots showing example performance a cuboid detector.

[0015] FIG. 8 is a flow diagram of an example process of training a cuboid detector.

[0016] FIG. 9 is a flow diagram of an example process of using a cuboid detector for cuboid detection and keypoint localization.

[0017] FIG. 10 schematically illustrates an example of a wearable display system, which can implement an embodiment of the deep cuboid detector.

[0018] Throughout the drawings, reference numbers may be re-used to indicate correspondence between referenced elements. The drawings are provided to illustrate example embodiments described herein and are not intended to limit the scope of the disclosure.

DETAILED DESCRIPTION

Overview

[0019] Models representing data relationships and patterns, such as functions, algorithms, systems, and the like, may accept input, and produce output that corresponds to the input in some way. For example, a model may be implemented as a machine learning method such as a convolutional neural network (CNN) or a deep neural network (DNN). Deep learning is part of a broader family of machine learning methods based on the idea of learning data representations as opposed to task specific algorithms and shows a great deal of promise in solving audio-visual computational problems useful for augmented reality, mixed reality, virtual reality, and machines intelligence. In machine learning, a convolutional neural network (CNN, or ConvNet) can include a class of deep, feed-forward artificial neural

networks, and CNNs have successfully been applied to analyzing visual imagery. Machine learning methods include a family of methods that can enable robust and accurate solutions to a wide variety of problems, including eye image segmentation and eye tracking.

[0020] Disclosed herein are examples of a cuboid detector which processes an input image of a scene and localizes at least one cuboid in the image. For example, a cuboid detector (such as a deep cuboid detector) can process a consumer-quality Red-Green-Blue (RGB) image of a cluttered scene and localize some or all three-dimensional (3D) cuboids in the image. A cuboid can comprise a boxy or a box-like object and can include a polyhedron (which may be convex) with, e.g., 4, 5, 6, 7, 8, 10, 12, or more faces. For example, cuboids can include pyramids, cubes, prisms, parallelepipeds, etc. Cuboids are not limited to such polyhedral shapes from geometry and can include box-like structures such as, e.g., appliances (e.g., television sets, computer monitors, toasters, washing machines, refrigerators), furniture (e.g., sofas, chairs, beds, cribs, tables, book cases, cabinets), vehicles (e.g., automobiles, buses), etc. As further described below, cuboids may be identified in terms of their faces, vertices, edges, or presence within a bounding box.

[0021] In some embodiments, a cuboid can comprise a geometric shape characterized as a tuple of N parameters. The parameters may be geometric in nature, like the radius of a sphere or the length, width, and height of the cuboid. A more general way to parameterize any geometric primitive can be to represent it as a collection of points on the surface of the primitive. If a random point on the surface of the primitive is chosen, the random point might not be localizable from a computer-vision point of view. It may be advantageous for the set of parameterization points to be geometrically informative and visually discriminative. For example, in the case of cuboids, the set of parameterization points may be the cuboid's vertices (which may be referred to sometimes herein as corners or keypoints).

[0022] In some embodiments, a cuboid is represented as a tuple of eight vertices, where each vertex can be denoted by its coordinates (e.g., Cartesian x,y coordinates) in the image. In such a representation, a cuboid is represented by 16 parameters: the two coordinates of each of the eight vertices. Not all 16 parameters might be needed in some cases, for example, as will be discussed below alternate cuboid representations may not

include some vertices (e.g., use only six vertices) and determine the other vertices using vanishing points.

[0023] Contrary to other approaches which fit a 3D model from low-level cues like corners, edges, and vanishing points, the cuboid detector disclosed herein can be an end-to-end deep learning system that detects cuboids across many semantic categories (e.g., ovens, shipping boxes, and furniture). In some implementations, the cuboid detector can localize a cuboid with a two-dimensional (2D) bounding box, and simultaneously localize the cuboid's keypoints (e.g., vertices or corners), effectively producing a 3D interpretation or representation of a box-like object. The cuboid detector can refine keypoints by pooling convolutional features iteratively, improving the accuracy of the keypoints detected. Based on an end-to-end deep learning framework, an advantage of some implementations of the cuboid detector is that there is little or no need to design custom low-level detectors for line segments, vanishing points, junctions, etc.

[0024] The cuboid detector can include a plurality of convolutional layers and non-convolutional layers of a convolutional neural network, a region proposal network (RPN), and a plurality of pooling and regressor layers. The RPN can generate object proposals in an image. The plurality convolutional layers and non-convolutional layers can generate a convolutional feature map of an input image. A convolutional layer of the CNN can include a kernel stack of kernels. A kernel of a convolutional layer, when applied to its input, can produce a resulting output activation map showing the response to that particular learned kernel. The resulting output activation map can then be processed by another layer of the CNN. Non-convolutional layers of the CNN can include, for example, a normalization layer, a rectified linear layer, or a pooling layer.

[0025] The region proposal network (RPN), which can be convolutional neural network or a deep neural network, can determine a 2D bounding box around a cuboid in the image from the convolutional feature map. The 2D bounding box can represent a region of interest (RoI) on the image which includes a cuboid at an image location. The plurality of pooling and regressor layers can include, for example, a pooling layer and two or more fully-connected layers (such as 3, 5, 10, or more layers). Based on the initial 2D bounding box, the plurality of cuboid pooling and regressor layers can, iteratively, determine a refined 2D bounding box and the cuboid's keypoints.

[0026] The cuboid detector can be trained in an end-to-end fashion and can be suitable for real-time applications in augmented reality (AR), mixed reality (MR), or robotics in some implementations. As described below, a wearable mixed reality display device (e.g., the wearable display system 1000 described with reference to FIG. 10) can include a processor programmed to perform cuboid detection on images acquired by an outward-facing camera of the display device. Some or all parameters of the cuboid detector can be learned in a process referred to as training. For example, a machine learning model can be trained using training data that includes input data and the correct or preferred output of the model for the corresponding input data. The machine learning model can repeatedly process the input data, and the parameters (e.g., the weight values) of the machine learning model can be modified in what amounts to a trial-and-error process until the model produces (or “converges” on) the correct or preferred output. For example, the modification of weight values may be performed through a process referred to as “back propagation.” Back propagation includes determining the difference between the expected model output and the obtained model output, and then determining how to modify the values of some or all parameters of the model to reduce the difference between the expected model output and the obtained model output.

Example Comparison of Object Detection and Cuboid Detection

[0027] Building a 3D representation of the world from a single monocular image is an important problem in computer vision. In some applications, objects having explicit 3D models are localized with their poses estimated. But without such 3D models, a person or a computer system (e.g., the wearable display system 1000 described with reference to FIG. 10) may still need to reason about its surrounding in terms of simple combinations of geometric shapes like cuboids, cylinders, and spheres. Such primitives, sometimes referred to as geons, can be easy for humans to reason about. Humans can effortlessly make coarse estimates about the pose of these simple geometric primitives and even compare geometric parameters like length, radius or area across disparate instances. While many objects are composed of multiple geometric primitives, a large number of real objects can be well approximated by as little as one primitive.

[0028] For example, a common shape is the box. Many everyday objects can geometrically be classified as a box (e.g., shipping boxes, cabinets, washing machines, dice, microwaves, desktop computers). Boxes (which are examples of cuboids) span a diverse set of everyday object instances, and humans can easily fit imaginary cuboids to these objects and localizing their vertices and faces. People can also compare the dimensions of different box-like objects even though they are not aware of the exact dimensions of the box-like objects or even if the objects are not perfect cuboids. Disclosed herein are systems and methods that implement a cuboid detector for detecting class agnostic geometric entities, such as cuboids. Class agnostic means that different classes of a geometric entity are not differentiated. For example, a cuboid detector may not differentiate between different classes of a cuboid, such as a shipping box, a microwave oven, or a cabinet. All of these box-like objects can be represented with the same simplified concept, a cuboid.

[0029] An embodiment of a cuboid detector can be used for 3D object detection as follows: fit a 3D bounding box to objects in an image (e.g., an RGB image or an RGB-Depth (RGB-D) image), detect 3D keypoints in the image, or perform 3D model to 2D image alignment. Because an image might contain multiple cuboids as well as lots of clutter (e.g., non-cuboidal objects), the cuboid detector can first determine a shortlist of regions of interest (RoIs) that correspond to cuboids. In addition to the 2D bounding box enclosing each cuboid, the cuboid detector can determine the location of all eight vertices.

[0030] Deep learning has revolutionized image recognition in the past few years. Many state-of-the-art methods in object detection today are built on top of deep networks that have been trained for the task for image classification. A cuboid detector can be a deep cuboid detector implementing one or more deep learning methods. The cuboid detector can have high accuracy and run in real-time using the hardware of a mobile device (e.g., the wearable display system 1000 described with reference to FIG. 10).

[0031] FIG. 1A is an example monocular image 100a illustrating two-dimensional (2D) object detection with a bounding box 104 overlaid around an object detected. FIG. 1B is an example monocular image 100b illustrating three-dimensional (3D) cuboid detection with a representation 108 of the cuboid overlaid on the object detected. FIG. 1B shows that one cuboid 108 inside the monocular image 100 is detected and its vertices localized. The eight vertices are shown as four black circles 112a-112d that are

connected by four edges 120a-120d (represented as dotted lines) and four additional black circles 116a-116d connected by four edges 124a-124d (represented as solid lines). Four of the vertices 112a-112d represent one face 128a of the cuboid, and the other four of the vertices 116a-116d represent another face 128b of the cuboid. The two faces 128a, 128b of the cuboid 108 are connected by four edges 132a-132d (represented as dashed lines) through the vertices 112a-112d, 116a-116d. The cuboid detector can detect box-like objects in a scene. Unlike object detection, the cuboid detector can determine more than a bounding box of an object. In addition, the cuboid detector can localize the vertices of the cuboids (e.g., compare FIG. 1A with FIG. 1B). In some embodiments, the cuboid detector can be class agnostic. For example, the cuboid detector does not care about the class of the cuboids being detected. For example, the cuboid detector can distinguish two classes of objects: a cuboid and a non-cuboid cuboid. The cuboid detector can perform 3D cuboid detection by determining all cuboids inside a monocular image and localize their vertices. The cuboid detector can be trained in an end-to-end fashion. The cuboid detector can run in real-time and perform cuboid detection with RGB images of cluttered scenes captured using a consumer-grade camera as input. A wearable display device (e.g., the wearable display system 1000 described with reference to FIG. 10) can implement the cuboid detector and use information about the detected cuboids to generate or update a world map indicative of the environment surrounding the user of the wearable display device.

[0032] A cuboid is a geometric object that can be parameterized, and a cuboid detector (e.g., a deep cuboid detector) can determine parameters of a cuboid in a scene. One approach to detect a cuboid is to detect the edges and try to fit the model of a cuboid to these edges. Hence, robust edge selection may be a useful aspect of the system. However, this becomes challenging when there are misleading textures on cuboidal surfaces, for example, if edges and corners are occluded or the scene contains considerable background clutter. It can be challenging to classify whether a given line belongs to a given cuboid with purely local features. The cuboid detector can learn to detect cuboids in images using a data-driven approach. The cuboid detector can assign a single label (e.g., “cuboid”) to box-like objects in a scene, even though the label is spread over many categories like houses, washing machines, ballot boxes, desks, cars, television sets, etc. The cuboid detector can include a CNN that is able to successfully learn features that help a system implementing it (e.g., the

wearable display system 1000 described with reference to FIG. 10) identify cuboids in different scenes.

[0033] In some embodiments, a cuboid detector can implement a deep learning model that jointly performs cuboid detection and keypoint localization. For example, a cuboid detector can include a deep neural network that jointly performs cuboid detection and keypoint localization. The cuboid detector can exceed the accuracy of the detection and localization accuracy performed by other methods. In some implementations, the cuboid detector can first detect the object of interest and then make coarse or initial predictions regarding the location of its vertices. The cuboid can utilize the coarse or initial predictions as an attention mechanism, performing refinement of vertices by only looking at regions with high probability of being a cuboid. In some embodiments, the cuboid detector can implement an iterative feature pooling mechanism to improve accuracy. The cuboid detector can combine cuboid-related losses and or implement alternate parametrizations to improve accuracy.

Example Cuboid Network Architecture and Loss Function

[0034] FIG. 2 depicts an example architecture of a cuboid detector. The cuboid detector 200 can include one or more of the following components: a convolutional layers 204 (also referred to herein as a CNN tower), a Region Proposal Network (RPN) 208, at least one pooling layer 212, or one or more fully connected layers 216 (e.g., a regional CNN (R-CNN) regressor (or classifier)). The pooling layer 212 and the fully connected layers 216 can implement iterative feature pooling, which refines cuboid keypoint locations. The R-CNN can be a Faster R-CNN.

[0035] The cuboid detector 200 can implement a deep cuboid detection pipeline. The first action of the deep cuboid detection pipeline can be determining Regions of Interest (RoIs) 220a1, 220b, in an image 202a where a cuboid might be present. The Region Proposal Network (RPN) 200 can be trained to output such RoIs 220a1, 220b as illustrated in the image 202b. Then, regions 224a with features corresponding to each RoI 220a1, 220b can be pooled, using one or more pooling layers 212, from a convolutional feature map 228 (e.g., the fifth convolutional feature map, conv5, in VGG-M from the Visual Geometry Group at Oxford University). These pooled features can be passed through two fully

connected layers 216. In some implementations, instead of just producing a 2D bounding box, the cuboid detector 200 can output the normalized offsets of the vertices from the center of the RoI 220a1, 220b. The cuboid detector 200 can refine the predictions by performing iterative feature pooling. The dashed lines in FIG. 2 show the regions 224a, 224b of the convolutional feature map 228, corresponding to the RoI 220a1 in the image 202b and a refined RoI 220a2 in the image 202c, from which features can be pooled. The two fully connected layers 216 can process the region 224b of the convolutional feature map 228 corresponding to the refined RoI 220a2 to determine a further refined RoI and/or a representation of a cuboid 232 in the image 202d.

[0036] The CNN Tower 204 can be the pre-trained fully convolutional part of ConvNets, such as VGG and ResNets. The convolutional feature map 228 refers to the output of the last layer of the CNN Tower 204. For example, the convolutional feature map 228 can be the output of the fifth convolutional layer, such as conv5 in VGG16 from the Visual Geometry Group at Oxford University with size $m \times n \times 512$).

[0037] The RPN 208 can be a fully convolutional network that maps every cell in the convolutional feature map 228 to a distribution over K multi-scale anchor-boxes, bounding box offsets, and objectness scores. The RPN can have two associated loss functions: a log loss function for objectness and a smooth L1 loss function for bounding box regression. The RPN 208 can, for example, use $512 \ 3 \times 3$ filters, then $18 \ 1 \times 1$ filters for objectness and $36 \ 1 \times 1$ filters for bounding box offsets.

[0038] The RoI pooling layer 212 can use, for example, max pooling to convert the features inside any valid region of interest 220a1, 220a2, 220b into a small fixed-size feature map (or a submap of the convolutional feature map 228). For example, for conv5 of size $m \times n \times 512$, the pooling layer 212 can produce an output of size $7 \times 7 \times 512$, independent of the input regions aspect ratio and scale. In some embodiments, spatial pyramid matching can be implemented.

[0039] The fully connected layers 216 (e.g., a R-CNN regressor) can then be applied to each fixed-size feature vector, outputting a cuboidness score, bounding box offsets (four numbers), and eight cuboid keypoint locations (16 numbers). The bounding box regression values (Δx , Δy , Δw , Δh) can be used to fit the initial object proposal tightly around the object. The keypoint locations can be encoded as offsets from the center of the RoI and

can be normalized by the proposal width/height as shown in FIG. 3. FIG. 3 illustrates RoI-normalized coordinates of vertices represented as offsets from the center of an RoI 304 in an image 300 and normalized by the region's width w and height h with (x_v, y_v) being a keypoint 308 and (x_c, y_c) being the center 312 of the RoI. Example ground truth targets for each keypoint are shown in Equations [1] and [2]:

$$t_x = \frac{x_v - x_c}{w}, \text{ and} \quad \text{Equation [1]}$$

$$t_y = \frac{y_v - y_c}{h}. \quad \text{Equation [2]}$$

[0040] Referring to FIG. 2, the R-CNN can include two fully connected layers 216 (e.g., 4096 neurons each) and can have three associated loss functions: a log loss function for cuboidness and smooth L1 loss functions for both bounding box and vertex regression.

[0041] When viewed in unison, the RoI pooling layer 212 and R-CNN layers act as a refinement mechanism, mapping an input box to an improved one, given the feature map. The cuboid detector 200 can apply the last part of the network multiple times (e.g., 2, 3, 4, or more times), referred to herein as iterative feature pooling.

[0042] The loss functions used in the RPN 208 can include $L_{anchor-cls}$, the log loss over two classes (e.g., cuboid vs. not cuboid) and $L_{anchor-reg}$, the Smooth L1 loss of the bounding box regression values for each anchor box. The loss functions for the R-CNN can include $L_{ROI-cls}$, the log loss over two classes (e.g., cuboid vs. not cuboid), $L_{ROI-reg}$, the Smooth L1 loss of the bounding box regression values for the RoI, and $L_{ROI-corner}$, the Smooth L1 loss over the RoI's predicted keypoint locations. The last term can be referred to as the corner or vertex regression loss. The complete loss function can be a weighted sum of the above mentioned losses and can be written as shown in Equation [3]. The loss weight λ_i can be different in different implementations, such as 0.1, 0.5, 1, 2, 5, 10, or more.

$$L = \lambda_1 L_{\text{anchor-cls}} + \lambda_2 L_{\text{anchor-reg}} + \lambda_3 L_{\text{ROI-cls}} + \lambda_4 L_{\text{ROI-reg}} + \lambda_5 L_{\text{ROI-corner}}. \quad \text{Equation [3]}$$

Example Performance

[0043] To determine its performance, an embodiment of the cuboid detector 200 was implemented using Caffe and built on top of an implementation of Faster R-CNN. To determine the performance, the VGG-M or VGG16 networks that have been pre-trained for the task of image classification on ImageNet were used. VGG-M is a smaller model with 7 layers while VGG16 contains 16 layers. All models were fine-tuned for 50K iterations using stochastic gradient descent (SGD) with a learning rate of 0.001, which was reduced by a factor of 10 after 30K iterations. Additional parameters used include a momentum of 0.9, weight decay of 0.0005, and dropout of 0.5. Instead of stage-wise training. Components of the cuboid detector 200 were jointly optimized with the values of all the loss weights as one (e.g., $\lambda_i = 1$ in Equation [3]).

[0044] Data. The SUN Primitive dataset (a comprehensive collection of annotated images covering a large variety of environmental scenes, places and the objects within; available from <https://groups.csail.mit.edu/vision/SUN/>) was used to train the deep cuboid detector 200. The dataset consists of 3516 images and is a mix of in-door scenes with lots of clutter, internet images containing only a single cuboid, and outdoor images of buildings that also look like cuboids. Both cuboid bounding boxes and cuboid keypoints have ground-truth annotations. This dataset includes 1269 annotated cuboids in 785 images. The rest of the images are negatives, e.g., they do not contain any cuboids. The dataset was split to create a training set of 3000 images and their horizontally flipped versions and a test set with 516 test images.

[0045] The cuboid detector 200 was evaluated on two tasks: cuboid bounding box detection and cuboid keypoint localization. For detection, a bounding box was correct if the intersection over union (IoU) overlap was greater than 0.5.2. Detections were sorted by confidence (e.g., the network's classifier softmax output) with the mean Average Precision (AP) as well as the entire Precision-Recall curve reported. For keypoint localization, the Probability of Correct Keypoint (PCK) and Average Precision of Keypoint (APK) metrics

were used to determine the cuboid detector's performance. PCK and APK are used in the human pose estimation literature to measure the performance of systems predicting the location of human body parts like head, wrist, etc. PCK measures the fraction of annotated instances that are correct when all the ground truth boxes are given as input to the system. A predicted keypoint was considered correct if its normalized distance from the annotation was less than a threshold (α). APK, on the other hand, takes both detection confidence and keypoint localization into consideration. A normalized distance, α , of 0.1 was used, meaning that a predicted keypoint was considered to be correct if it lied within a number of pixels of the ground truth annotation of the keypoint shown in Equation [4]. The normalized distance, α , can be different in different implementations, such as 0.01, 0.2, 0.3, 0.5, 0.9, or more.

$$1.1 * \max(\text{height}, \text{width}) \quad \text{Equation [4]}$$

See FIGS. 7A-7F for these metrics reported on the SUN Primitive test set and samples of cuboid detections and vertices localization in monocular images 400a-400y, 404a-404e illustrated in FIGS. 4A-4G. For example, FIG 4A shows a monocular image 400a with example representations 108a-108d of four cuboids each represented as eight vertices. As another example, FIG 4A shows another monocular image 400b with an example representation 108a of a cuboid with four vertices representing one face of the cuboid connected by four edges (shown as solid lines) and four vertices representing another face of the cuboid connected by another four edges (shown as dotted lines). The eight vertices on these two faces of the representation 108a of the cuboid are connected by four edges (shown as dashed lines).

[0046] FIGS. 7A-7F are graphs illustrating example deep cuboid detector evaluation metrics. APK: Average Precision of Keypoint, PCK: Probability of Correct Keypoint: Normalized distance from GT corners, Order of keypoints: front-top-left, back-top-left, front-bottom-left, front-top-right, back-bottom-left, front-bottom-right, back-top-right, back-bottom-right. B: bounding box loss, C: corner loss, and I: iterative. FIGS. 4A-4F show images illustrating example cuboid detection and keypoint location using VGG16 as the CNN tower and iterative feature pooling. The cuboid detector 200 was able to localize the vertices of cuboids in consumer-grade RGB images. The cuboid detector 200 was able to handle both objects like boxes (that are perfectly modeled by a cuboid) as well as objects like sinks (that are only approximate cuboids). FIG. 4G show example images 404a-404e

illustrating improper cuboid detection and keypoint localization, which can be reduced or eliminated as further described below.

[0047] In one implementation, the cuboid detector 2 achieved a mAP of 75.47 for bounding box detection, which was significantly better than the HOG-based system with a mAP of 24.0.

[0048] Multi-Task learning. Multiple network each performing different multiple tasks were trained. A base network that just output bounding boxes around cuboids was trained. This base network performed general object detection using rectangles enclosing cuboids. The base network output the class of the box and the bounding box regression values. Next, a different network with additional supervision about the location of the corners was trained. This network did not output bounding box regression coordinates. Then, a network (e.g., the cuboid detector 200) that output both the bounding box regression values and the coordinates of the vertex was trained. A corresponding term was added to the loss function for each additional task. From testing, adding more tasks (bounding box detection, keypoint localization, or both bounding box detection and keypoint localization), affected the performance of the cuboid detector (see Table 1).

Table 1. Multi-task learning Results. A network was trained using only the bounding box loss, then using the cuboid corner loss.

Additional loss function	AP	APK	PCK
Bounding Box Loss	66.33	-	-
Corner Loss	58.39	28.68	27.64
Bounding Box + Corner Loss	67.11	34.62	29.38

[0049] Iterative Feature Pooling. In R-CNN, the final output is a classification score and the bounding box regression values for every region proposal. The bounding box regression allows moving the region proposal around and scaling it such that the final bounding box localizes just the object. This implies that the initial region from which the features are pooled to make this prediction was not entirely correct. In some embodiments, the cuboid detector 200 goes back and pools features from the refined bounding box. This can be implemented in the network itself, meaning that the cuboid detector 200 performs

iterative bounding box regression while training and testing in exactly the same way. The input to the fully-connected layers 216 of the regressor is a fixed-size feature map, a submap the convolutional feature map 228, that includes of the pooled features from different region proposals from conv5 layer. The R-CNN outputs can be used for bounding box regression on the input object proposals to produce new proposals. Then features can be pooled from these new proposals and passed through the fully-connected layers 216 of the regressor again. In some embodiments, the cuboid detector 200 is an “any-time prediction system” where for applications which are not bound by latency, bounding box regression can be performed more than once. The performance results (see Table 2) show that iterative feature pooling can greatly improve both bounding box detection and vertex localization (see FIGS. 5A-5C). There was not a significant change in performance when features were iteratively pooled two or more times (e.g., 2, 3, 4, 5, 6, or more times). In some implementations, two iterations are used. FIGS. 5A-5C show example images 500a1-500i1, 500a2-500i2 illustrating improved performance (e.g., compare the representations 108b1, 108b2 of the cuboid in images 500a1, 500a2 and the shape of the bookcase 504 in these images 504. with keypoint refinement via iterative feature pooling. Cuboid detection regions were refined by re-pooling features from conv5 using the predicted bounding boxes.

Table 2. Results for Iterative Feature Pooling. Iterative feature pooling improved the box detection AP by over 4% and PCK over 7%.

Method	AP	APK	PCK
Corner Loss	58.39	28.68	27.64
Corner Loss + Iterative	62.89	33.98	35.56
BB + Corner Losses	67.11	34.62	29.38
BB + Corner Loss + Iterative	71.72	37.61	36.53

[0050] Depth of Network. Two base models, VGG16 and VGG-M, were tested. While VGG16 has a very deep architecture with 16 layers, VGG-M is a smaller model with 7 layers. Table 3 shows the results of the testing. Interestingly, for this dataset and task, two iterations through the shallower network outperformed one iteration through the deeper network. Coupled with the fact the shallower network with iteration run twice as fast, a

cuboid detector 200 can advantageously include a shallower CNN tower with fewer than 10 layers (e.g., 5, 7, or 9 layers). In some embodiments, a cuboid detector 200 can include a deeper CNN tower (e.g., 12, 15, 20, or more layers). The four model tested each had average precision (AP) higher than the AP of a HOG-based system (24.0).

Table 3. VGG-M (7 layers) vs. VGG16 (16 layers) base network. I: iterative feature pooling was performed. The deeper cuboid detector outperformed the shallower one.

Method	AP	APK	PCK	Size	Speed
VGG-M	67.11	34.62	29	334 MB	14 fps
VGG-M + I	71.72	37.61	36	334 MB	10 fps
VGG16	70.50	33.65	35	522 MB	5 fps
VGG16 + I	75.47	41.21	38	522 MB	4 fps

[0051] Effect of Training Set Size. The impact of increasing the size of training data was measured. Three datasets of varying sizes, 1K, 2K and 3K images, were created and used to train a common network (VGG-M + Iterative). The results (see Table 4) show significantly improved performance when using larger training set sizes.

Table 4. Performance vs. number of training images. Deep cuboid detection can benefit from more training images.

Number of Images	AP	APK	PCK
1000	40.47	20.83	26.60
2000	52.17	27.51	29.31
3000	71.72	37.61	26.53

[0052] Memory and Runtime Complexity. The cuboid detector 200 was able to run at interactive rates on a Titan Z GPU while the HOG-based approach would take minutes to process a single image. The real-time nature of the system may be the result of Faster R-CNN being used as the regressor. In some embodiments, the cuboid detector 200 can implement a single show multibox detector (SSD) to further improve its speed performance.

Table 3 shows the model sizes, which can be reduced to on mobile devices (e.g., the wearable display system 1000 described with reference to FIG. 10).

Example Keypoint Parameterizations

[0053] An embodiment of the cuboid detector 200 can output a cuboid's vertices directly. Many convex cuboids have eight vertices, six faces, and twelve edges (not all of which may be visible in an image). However, certain viewpoints may have an inherent ambiguity, which may have led to the improper cuboid identification shown in FIG. 4G. For example, which face of the cube in FIG. 4G should be labelled the front? Since the cuboid detector 200 detector may need to deal with such configurations, alternate cuboid parametrizations were explored. If the world origin is considered to coincide with camera center coordinates, a parameterization of a cuboid can be represented with 12 numbers. The following parameterization may be minimal; in other parameterizations, additional or different parameters can be used.

(X, Y, Z) – Coordinates of the center of the cuboid in 3D

(L, W, H) - Dimensions of the cuboid

(θ, ψ, ϕ) - 3 angles of rotation of the cuboid (e.g., Euler angles)

(f, C_x, C_y) - Intrinsic camera parameters (e.g., focal length and coordinates of the optical center)

[0054] For many modern cameras, no skew in the camera and equal focal lengths (in orthogonal directions) can be assumed. The over-parameterization of a cuboid (e.g., a sixteen-parameter parameterization of a cuboid) may allow a cuboid detector 200 to produce outputs that do not represent cuboids (see, e.g., some examples in FIG. 4G). Several different re-parameterizations of a cuboid were tested to better utilize the geometric constraints. In general, the test results show that the network was able to learn features for tasks that had more visual evidence in the image and predict parameters which can be scaled properly for stable optimization. When dealing with 3D geometry and deep learning, proper parametrization is advantageous. Even image-to-image transformations, such as like homographies (e.g., isomorphisms of projected spaces) may benefit from re-parametrization

(e.g., the four-point parametrization). Such techniques may reduce or eliminate improper identification of cuboids in images.

[0055] Six-corner parametrization. An alternate parameterization in which only six coordinates of eight cuboid vertices were predicted by the detector. The locations of the remaining two coordinates were inferred using the relationship that there may be parallel edges in cuboids. For example, the edges that are parallel in 3D meet at the vanishing point in the image. There may be two pairs of parallel lines on the top base of the cuboid 600 and two pairs of parallel lines on the bottom face of the cuboid. The pair of parallel lines 604a, 604b on the top face of the cuboid 600 and the pair parallel line 606a, 606b on the bottom face of the cuboid should meet at the same vanishing point 608a as shown in FIG. 6. The pair of parallel lines 604c, 604d on the top face of the cuboid 600 and the pair parallel line 606c, 606d on the bottom face of the cuboid should meet at the same vanishing point 608b. Accordingly, the position of the remaining two points 612a, 612b can be inferred. This allows a cuboid detector 200 to parameterize an output of 12 numbers in some implementations. FIG. 6 schematically illustrates example cuboid vanishing points 608a, 608b. Vanishing points 608a, 608b produced by extrapolating the edges of a cube form a vanishing line 616 and can be used to reduce the number of parameters. The Front-Top-Left (FTL) keypoint 612a and Back-Bottom-Right (BBR) keypoint 612b can be excluded from the parametrization and inferred using estimated vanishing points (VPs) techniques.

[0056] Eight-corner parameterization was compared with six-corner parameterization. The ground truth data for two vertices was not used while training. One vertex from each the back and front faces was dropped (those whose detection rates (PCK) were the worst). A network was trained to predict the location of the remaining six corners. The locations of the two dropped vertices were inferred using these six corners. The cuboid detector 200 first determined the vanishing points corresponding to the six points predicted. This re-parameterization may lead to a reduction in performance (see Table 5). This degradation may be due to the fact that visual evidence corresponding to the two inferred corners present in the image was discarded. Also, any error in prediction of one vertex due to occlusion or any other reason would directly propagate to the inferred corners. However, left to the cuboid detector 200, it learned multiple models to detect a cuboid. The network of the cuboid detector 200 was free to use all visual evidence to localize the corners of the

cuboid. The cuboid detector 200 was capable of doing pure geometric reasoning because in many cases the corner on the back did not have visual evidence in the image due to self-occlusion.

Table 5. Eight-Corner vs. six-corner parameterization. Eight-corner parameterization uses all of the cuboid's corners, whereas in the six-corner parameterization, the BBR and FTL corners are dropped (see FIG. 6) and inferred from the vanishing points. This shows how an example network was able to do geometric reasoning and the over-parameterization may add robustness to the system. BBR: Back-Bottom-Right and FTL: Front-Top-Left.

Method	AP	APK	PCK	PCK of BBR Corner	PCK of FTL Corner	PCK of Remaining Corners
6 corners	65.26	29.64	27.36	24.44	21.11	28.89
8 corners	67.11	34.62	29.38	27.22	29.44	29.73

[0057] Vanishing point parametrization: Another re-parameterization uses locations of the two vanishing points and the slopes of six lines which will form the edges of the cuboid (see FIG. 6). Note that these vanishing points correspond to a particular cuboid and might be different from the vanishing point of the entire image. The intersection points of these six lines would give the vertices of the cuboid in this example. However, the locations of the vanishing points many lie outside the region of interest and have little or confounding visual evidence in the region of interest or the entire image itself. It also may become difficult to normalize the targets to predict the vanishing points directly. The slopes of the six lines can vary between $-\infty$ and $+\infty$. Instead of predicting the slope directly, the slopes can be regressed to the value of $\sin(\tan^{-1}(\theta))$. There can exist a set of hyperparameters (e.g., loss weights, learning rates, solver, etc.) for which an embodiment of this network can be trained.

Example Process of Training a Cuboid Detector

[0058] FIG. 8 is a flow diagram of an example process 800 of training a cuboid detector. The process 800 starts at block 804, where a plurality of training images each

comprising at least one cuboid is received. Some of the training images can each include one or more cuboids. The process 800 can include performing a cuboid-specific (e.g., cuboid-specific) data augmentation strategy to improve the performance of a trained cuboid detector. At block 808, a convolutional neural network is received. The convolutional neural network can be trained for objection detection. For example, the convolutional neural network can be VGG16 or VGG-M. The convolutional neural network can be a deep neural network in some implementations.

[0059] At block 812, a cuboid detector is generated. The cuboid detector can include a CNN tower. The CNN tower can include a plurality of convolutional layers and non-convolutional layers of the convolutional neural network received at block 808. For example, the CNN tower can include some or all convolutional layers of the convolutional neural network received. The non-convolutional layers can include a normalization layer, a brightness normalization layer, a batch normalization layer, a rectified linear layer, an upsampling layer, a concatenation layer, a pooling layer, a softsign layer, or any combination thereof. The CNN tower can generate a convolutional feature map from an input image, such as a monocular image.

[0060] The cuboid detector can include a region proposal network (RPN), such as a CNN or a DNN. The region proposal network can be connected to a layer of the CNN tower. The region proposal network can determine a region of interest (RoI) comprising a cuboid in the image using the convolutional feature map. For example, the region of interest can be represented as a two-dimensional (2D) bounding box enclosing a cuboid at a cuboid image location. The cuboid can comprise a cuboid, a cylinder, a sphere, or any combination thereof. The RPN can be associated with at least two loss functions, such as a log loss function and a smooth L1 loss function during training.

[0061] The cuboid detector can include a pooling layer and at least one regressor layer. The pooling layer can be connected to a layer of the CNN tower. The pooling layer can determine, using the cuboid image location, a submap of the convolutional feature map corresponding to the region of interest comprising the cuboid. The pooling layer and the region proposal network can be connected to the same layer of the CNN tower.

[0062] The cuboid detector can include two regressor layers, such as two fully-connected layers, of a regional-CNN (R-CNN) or a fast R-CNN. As another example, the

regressor layer is not fully connected. The regressor layer can be associated with at least three loss functions during training. For example, the at least three loss functions comprises a log loss function and a smooth L1 loss function.

[0063] The cuboid detector can be trained. At block 816, the cuboid detector can determine a region of interest at an image location comprising a cuboid in a training image received at block 804. In some embodiments, a representation of the cuboid in the image can be determined. To determine the RoI at the cuboid image location and the representation of the cuboid, the cuboid detector can generate a convolutional feature map for the training image using the convolutional layers and non-convolutional layers of the CNN tower. Based on the convolutional feature map, the region proposal network can determine the RoI comprising the cuboid at an initial image location in the training image. Based on the initial image location of the cuboid in the training image, the pooling layer of the cuboid detector can determine a submap of the convolutional feature map corresponding to the RoI comprising the cuboid at the initial image location. The at least one regression layer can determine the RoI at the cuboid image location and the representation of the cuboid. The initial cuboid image location or the cuboid image location can be represented as a two-dimensional (2D) bounding box. In some implementations, the method 800 can include iteratively determining, using the pooling layer, the at least one regressor layer, and the submap of the convolutional feature map corresponding to the RoI comprising the cuboid, the RoI at the cuboid image location and the representation of the cuboid.

[0064] The representation of the cuboid can be different in different implementations. The representation can include a parameterized representation of the cuboid. For example, the parameterized representation of the cuboid can include locations of a plurality of keypoints of the cuboid (e.g., a cuboid) in the image, such as six or eight vertices of the cuboid in the image. As another example, the parameterized representation can include normalized offsets of the plurality of keypoints of the cuboid from the center of the image. As a further example, the parameterized representation comprises N tuples, such as 6 tuples. As an example, the parameterized representation of the cuboid comprises a vanishing point parameterization.

[0065] At block 820, a first difference between a reference image location and the determined image location and a second difference between a reference representation of the

cuboid and the determined representation of the cuboid can be determined. The reference representation of the cuboid can include the ground truth targets for each keypoint as illustrated in Equations [1] and [2] above. The reference image location can include a bounding box represented by the ground truth targets.

[0066] At block 824, weights of the cuboid detector can be updated based on the first difference and the second difference. The differences can be represented as the loss function (or components thereof) shown in Equation [3]. Some or all of the weights of the cuboid detector can be updated based on the differences determined. For example, the weights of the region proposal network and the weights of the at least one regressor layer can be updated based on the differences. As another example, the weights of the RPN and the weights of the at least one regressor layer can be updated without updating the weights of the first CNN based on the differences. As a further example, the weights of the CNN tower, the weights of the region proposal network, and the weights of the at least one regressor layer can be updated based on the differences. The process 800 can optionally include training the cuboid detector from a larger dataset and synthetic data, network optimization, and regularization techniques to improve generalization.

Example Process of Using a cuboid detector for Cuboid Detection and Keypoint Localization

[0067] FIG. 9 is a flow diagram of an example process 900 of using a cuboid detector for cuboid detection and keypoint localization. The process 900 starts at block 904, where a system (e.g., the wearable display system 1000 described with reference to FIG. 10) receives an input image including a possible cuboid. The image can include one or more cuboids. The image can comprise a color image (e.g., RGB or RGB-D) and the image may be monocular. The image may be a frame of a video and may be obtained using the outward-facing imaging system 1044 of the wearable display system 1000 described with reference to FIG. 10.

[0068] At block 908, the wearable display system 1000 can access a cuboid detector (such as the cuboid detector trained by the process 800 illustrated in FIG. 8). The cuboid detector can include a CNN tower comprising a plurality of convolutional layers and non-convolutional layers. The cuboid detector can include a region proposal network

connected to the CNN tower. The cuboid detector can include a pooling layer and at least one regressor layer. The pooling layer can be connected to the CNN tower.

[0069] At block 912, the wearable display system 1000 can generate, using the plurality of convolutional layers and the non-convolutional layers of the CNN tower and the image, a convolutional feature map (e.g., the convolutional feature map 228 in FIG. 2). At block 916, the wearable display system 1000 can determine, using the region proposal network, at least one RoI comprising a cuboid at a cuboid image location of the image (e.g., the regions of interest 220a1, 220a2, 220b in FIG. 2). The cuboid image location can be represented as a two-dimensional (2D) bounding box. At block 920, the wearable display system 1000 can determine, using the pooling layer (e.g., the pooling layer 212 in FIG. 2) and the cuboid image location, a submap of the convolutional feature map corresponding to the region of interest comprising the cuboid. For example, the submap can be determined from the regions 224a of the convolutional feature map 228 from which the features can be pooled in FIG. 2. At block 924, the wearable display system 1000 can determine, using the regressor layer (e.g., a R-CNN regressor) and the submap, a refined RoI at a refined cuboid image location and a representation of the cuboid. The refined cuboid image location can be represented as a two-dimensional (2D) bounding box.

[0070] In some embodiments, the method 900 includes iterative feature pooling. For example, the wearable display system 1000 can determine using the refined cuboid image location, a refined submap of the convolutional feature map corresponding to the refined region of interest comprising the cuboid. For example, the submap can be determined from the regions 224b of the convolutional feature map 228 from which the features can be pooled in FIG. 2. The wearable display system 1000 can determine, using the pooling layer, the at least one regressor layer, and the refined submap of the convolutional feature map corresponding to the refined RoI, a further refined RoI at a further refined cuboid image location and a further defined representation of the cuboid.

[0071] The wearable display system 1000 can interact with a user of the system based on the refined region of interest at the refined cuboid image location and the representation of the cuboid. For example, the cuboid can correspond to a stationary box, and the wearable display system 1000 can generate character animation in relation to the stationary box based on the refined image location of the cuboid and the representation of the

cuboid. As another example, the cuboid can correspond to a hand-held cuboid. The wearable display system 1000 can determine a pose of the cuboid using the representation of the cuboid, and interact with the user of the system based on the pose of the cuboid. As a further example, the cuboid can correspond to a rare object not recognizable by a CNN. The wearable display system 1000 can provide the user with a notification that the rare object not recognizable by the CNN is detected. As an example, the cuboid corresponds to a man-made structure (e.g., a building). The wearable display system 1000 can assist the user of the system during an unmanned flight based on the refined RoI at the refined cuboid image location and the representation of the cuboid. As another example, the cuboid can be a cuboid that corresponds to a marker. The wearable display system 1000 can perform simultaneous location and mapping (SLAM) based on the pose of the cuboid.

Example Applications

[0072] Detecting box-like objects in images and extracting 3D information like pose can help overall scene understanding. Many high-level semantic problems can be tackled by first detecting boxes in a scene (e.g., extracting the free space in a room by reducing the objects in a scene to boxes, estimating the support surfaces in the scene and estimating the scene layout).

[0073] The cuboid detectors disclosed herein can open up one or more possibilities for augmented reality (AR), human-computer interaction (HCI), autonomous vehicles, drones, or robotics in general. For example, the cuboid detector can be used as follows.

[0074] For Augmented Reality, cuboid vertex localization followed by 6-degree of freedom (6-dof) pose estimation allows a content creator to use the cuboid-centric coordinate system defined by a stationary box to drive character animation. Because the volume of space occupied by the stationary cuboid is known based on cuboid vertex location followed by 6-dof pose estimation, animated characters can jump on the box, hide behind it, and even start drawing on one of the box's faces. Accordingly, a content creator can use the cuboid detector to build dynamic worlds around cuboids.

[0075] For Human-Computer Interaction, users may interact with scenes using boxy objects around them. A content creator may create a game or user environment in

which worlds are built up from cuboids. As another example, a hand-held cuboid can be used as a lightweight game controller. A system, such as the wearable display system 1000 described with reference to FIG. 10, can include a camera capturing images of the hand-held cube over time. And the system can estimate the cube's pose, effectively tracking the cube in 3D space, using the images captured. In some embodiments, the cuboid can serve as a way to improve interaction in AR systems (e.g., the tabletop AR demo using cuboids).

[0076] For autonomous vehicles, 3D cuboid detection allows the vehicle to reason about the spatial extent of rare objects that might be missing in supervised training set. By reasoning about the pose of objects in a class-agnostic manner, autonomous vehicles can be safer drivers.

[0077] For drones, man-made structures, such as buildings, houses, or cars, can be well-approximated with cuboids, assisting navigation during unmanned flights. For robotics in general, detecting box-like objects in images and extracting their 3D information like pose helps overall scene understanding. For example, placing a handful of cuboids in a scene (instead of Aruco markers) can make pose tracking more robust for simultaneous location and mapping (SLAM) applications.

Additional Embodiments

[0078] In some embodiments, the cuboid detector does not rely on bottom-up image processing and works satisfactorily on real images in real-time. The cuboid detector can be trained using a large training database of 3D models and some kind of learning for 2D-to-3D alignment. In some implementations, the cuboid detector can implement a geometry-based method, a deformable parts model, a histogram oriented gradients (HOG)-based model (e.g., a HOG classifier). The cuboid detector can detect cuboid vertices in different views and determine a final cuboid configuration based on a score from the HOG classifier, 2D vertex displacement, edge alignment score and a 3D shape score that takes into account how close the predicted vertices are to a cuboid in 3D. The cuboid detector can jointly optimize over visual evidence (corners and edges) found in the image while penalizing the predictions that stray too far from an actual 3D cuboid.

[0079] Without being limited by theory, the cuboid detector may owe its performance convolutional neural networks. A CNN can be superior to existing methods for

the task of image classification. To localize a cuboid in an image, the image is broken down into regions and these regions are classified instead, for example, in real-time. The cuboid detector can perform detection in a single step. A cuboid detector, for example, running on the wearable display system 1000 described with reference to FIG. 10, can process 50-60 frames per second, thus performing real-time cuboid detection and keypoint localization. The iterative keypoint refinement implemented by the cuboid detector can be based on iterative error feedback approach of, the network cascades in, the iterative bounding box regression of Multi-Region CNN and Inside-Outside Networks. Alternatively, or additionally, the iterative keypoint refinement implemented by the cuboid detector can be based on a Recurrent Neural Networks.

Example NN Layers

[0080] A layer of a neural network (NN), such as a deep neural network (DNN) can apply a linear or non-linear transformation to its input to generate its output. A deep neural network layer can be a normalization layer, a convolutional layer, a softsign layer, a rectified linear layer, a concatenation layer, a pooling layer, a recurrent layer, an inception-like layer, or any combination thereof. The normalization layer can normalize the brightness of its input to generate its output with, for example, L2 normalization. The normalization layer can, for example, normalize the brightness of a plurality of images with respect to one another at once to generate a plurality of normalized images as its output. Non-limiting examples of methods for normalizing brightness include local contrast normalization (LCN) or local response normalization (LRN). Local contrast normalization can normalize the contrast of an image non-linearly by normalizing local regions of the image on a per pixel basis to have a mean of zero and a variance of one (or other values of mean and variance). Local response normalization can normalize an image over local input regions to have a mean of zero and a variance of one (or other values of mean and variance). The normalization layer may speed up the training process.

[0081] The convolutional layer can apply a set of kernels that convolve its input to generate its output. The softsign layer can apply a softsign function to its input. The softsign function ($\text{softsign}(x)$) can be, for example, $(x / (1 + |x|))$. The softsign layer may neglect impact of per-element outliers. The rectified linear layer can be a rectified linear

layer unit (ReLU) or a parameterized rectified linear layer unit (PReLU). The ReLU layer can apply a ReLU function to its input to generate its output. The ReLU function $\text{ReLU}(x)$ can be, for example, $\max(0, x)$. The PReLU layer can apply a PReLU function to its input to generate its output. The PReLU function $\text{PReLU}(x)$ can be, for example, x if $x \geq 0$ and ax if $x < 0$, where a is a positive number. The concatenation layer can concatenate its input to generate its output. For example, the concatenation layer can concatenate four 5×5 images to generate one 20×20 image. The pooling layer can apply a pooling function which down samples its input to generate its output. For example, the pooling layer can down sample a 20×20 image into a 10×10 image. Non-limiting examples of the pooling function include maximum pooling, average pooling, or minimum pooling.

[0082] At a time point t , the recurrent layer can compute a hidden state $s(t)$, and a recurrent connection can provide the hidden state $s(t)$ at time t to the recurrent layer as an input at a subsequent time point $t+1$. The recurrent layer can compute its output at time $t+1$ based on the hidden state $s(t)$ at time t . For example, the recurrent layer can apply the softsign function to the hidden state $s(t)$ at time t to compute its output at time $t+1$. The hidden state of the recurrent layer at time $t+1$ has as its input the hidden state $s(t)$ of the recurrent layer at time t . The recurrent layer can compute the hidden state $s(t+1)$ by applying, for example, a ReLU function to its input. The inception-like layer can include one or more of the normalization layer, the convolutional layer, the softsign layer, the rectified linear layer such as the ReLU layer and the PReLU layer, the concatenation layer, the pooling layer, or any combination thereof.

[0083] The number of layers in the NN can be different in different implementations. For example, the number of layers in the DNN can be 50, 100, 200, or more. The input type of a deep neural network layer can be different in different implementations. For example, a layer can receive the outputs of a number of layers as its input. The input of a layer can include the outputs of five layers. As another example, the input of a layer can include 1% of the layers of the NN. The output of a layer can be the inputs of a number of layers. For example, the output of a layer can be used as the inputs of five layers. As another example, the output of a layer can be used as the inputs of 1% of the layers of the NN.

[0084] The input size or the output size of a layer can be quite large. The input size or the output size of a layer can be $n \times m$, where n denotes the width and m denotes the height of the input or the output. For example, n or m can be 11, 21, 31, or more. The channel sizes of the input or the output of a layer can be different in different implementations. For example, the channel size of the input or the output of a layer can be 4, 16, 32, 64, 128, or more. The kernel size of a layer can be different in different implementations. For example, the kernel size can be $n \times m$, where n denotes the width and m denotes the height of the kernel. For example, n or m can be 5, 7, 9, or more. The stride size of a layer can be different in different implementations. For example, the stride size of a deep neural network layer can be 3, 5, 7 or more.

[0085] In some embodiments, a NN can refer to a plurality of NNs that together compute an output of the NN. Different NNs of the plurality of NNs can be trained for different tasks. A processor (e.g., a processor of the local data processing module 1024 described with reference to FIG. 10) can compute outputs of NNs of the plurality of NNs to determine an output of the NN. For example, an output of a NN of the plurality of NNs can include a likelihood score. The processor can determine the output of the NN including the plurality of NNs based on the likelihood scores of the outputs of different NNs of the plurality of NNs.

Example Wearable Display System

[0086] In some embodiments, a user device can be, or can be included, in a wearable display device, which may advantageously provide a more immersive virtual reality (VR), augmented reality (AR), or mixed reality (MR) experience, where digitally reproduced images or portions thereof are presented to a wearer in a manner wherein they seem to be, or may be perceived as, real.

[0087] Without being limited by theory, it is believed that the human eye typically can interpret a finite number of depth planes to provide depth perception. Consequently, a highly believable simulation of perceived depth may be achieved by providing, to the eye, different presentations of an image corresponding to each of these limited number of depth planes. For example, displays containing a stack of waveguides may be configured to be worn positioned in front of the eyes of a user, or viewer. The stack

of waveguides may be utilized to provide three-dimensional perception to the eye/brain by using a plurality of waveguides to direct light from an image injection device (e.g., discrete displays or output ends of a multiplexed display which pipe image information via one or more optical fibers) to the viewer's eye at particular angles (and amounts of divergence) corresponding to the depth plane associated with a particular waveguide.

[0088] In some embodiments, two stacks of waveguides, one for each eye of a viewer, may be utilized to provide different images to each eye. As one example, an augmented reality scene may be such that a wearer of an AR technology sees a real-world park-like setting featuring people, trees, buildings in the background, and a concrete platform. In addition to these items, the wearer of the AR technology may also perceive that he "sees" a robot statue standing upon the real-world platform, and a cartoon-like avatar character flying by which seems to be a personification of a bumble bee, even though the robot statue and the bumble bee do not exist in the real world. The stack(s) of waveguides may be used to generate a light field corresponding to an input image and in some implementations, the wearable display comprises a wearable light field display. Examples of wearable display device and waveguide stacks for providing light field images are described in U.S. Patent Publication No. 2015/0016777, which is hereby incorporated by reference herein in its entirety for all it contains.

[0089] FIG. 10 illustrates an example of a wearable display system 1000 that can be used to present a VR, AR, or MR experience to a display system wearer or viewer 1004. The wearable display system 1000 may be programmed to perform any of the applications or embodiments described herein (e.g., executing CNNs, reordering values of input activation maps or kernels, eye image segmentation, or eye tracking). The display system 1000 includes a display 1008, and various mechanical and electronic modules and systems to support the functioning of that display 1008. The display 1008 may be coupled to a frame 1012, which is wearable by the display system wearer or viewer 1004 and which is configured to position the display 1008 in front of the eyes of the wearer 1004. The display 1008 may be a light field display. In some embodiments, a speaker 1016 is coupled to the frame 1012 and positioned adjacent the ear canal of the user in some embodiments, another speaker, not shown, is positioned adjacent the other ear canal of the user to provide for stereo/shapeable sound control. The display system 1000 can include an outward-facing

imaging system 1044 (e.g., one or more cameras) that can obtain images (e.g., still images or video) of the environment around the wearer 1004. Images obtained by the outward-facing imaging system 1044 can be analyzed by embodiments of the deep cuboid detector to detect and localize cuboids in the environment around the wearer 1004.

[0090] The display 1008 is operatively coupled 1020, such as by a wired lead or wireless connectivity, to a local data processing module 1024 which may be mounted in a variety of configurations, such as fixedly attached to the frame 1012, fixedly attached to a helmet or hat worn by the user, embedded in headphones, or otherwise removably attached to the user 1004 (e.g., in a backpack-style configuration, in a belt-coupling style configuration).

[0091] The local processing and data module 1024 may comprise a hardware processor, as well as non-transitory digital memory, such as non-volatile memory e.g., flash memory, both of which may be utilized to assist in the processing, caching, and storage of data. The data include data (a) captured from sensors (which may be, e.g., operatively coupled to the frame 1012 or otherwise attached to the wearer 1004), such as image capture devices (such as cameras), microphones, inertial measurement units, accelerometers, compasses, GPS units, radio devices, and/or gyros; and/or (b) acquired and/or processed using remote processing module 1028 and/or remote data repository 1032, possibly for passage to the display 1008 after such processing or retrieval. The local processing and data module 1024 may be operatively coupled to the remote processing module 1028 and remote data repository 1032 by communication links 1036, 1040, such as via a wired or wireless communication links, such that these remote modules 1028, 1032 are operatively coupled to each other and available as resources to the local processing and data module 1024. The image capture device(s) can be used to capture the eye images used in the eye image segmentation, or eye tracking procedures.

[0092] In some embodiments, the remote processing module 1028 may comprise one or more processors configured to analyze and process data and/or image information such as video information captured by an image capture device. The video data may be stored locally in the local processing and data module 1024 and/or in the remote data repository 1032. In some embodiments, the remote data repository 1032 may comprise a digital data storage facility, which may be available through the internet or other networking configuration in a “cloud” resource configuration. In some embodiments, all data is stored

and all computations are performed in the local processing and data module 1024, allowing fully autonomous use from a remote module.

[0093] In some implementations, the local processing and data module 1024 and/or the remote processing module 1028 are programmed to perform embodiments of reordering values of input activation maps or kernels, eye image segmentation, or eye tracking disclosed herein. For example, the local processing and data module 1024 and/or the remote processing module 1028 can be programmed to perform embodiments of the process 900 described with reference to FIG. 9. The local processing and data module 1024 and/or the remote processing module 1028 can be programmed to perform cuboid detection and keypoint localization disclosed herein. The image capture device can capture video for a particular application (e.g., augmented reality (AR), human-computer interaction (HCI), autonomous vehicles, drones, or robotics in general). The video can be analyzed using a CNN by one or both of the processing modules 1024, 1028. In some cases, off-loading at least some of the reordering values of input activation maps or kernels, eye image segmentation, or eye tracking to a remote processing module (e.g., in the “cloud”) may improve efficiency or speed of the computations. The parameters of the CNN (e.g., weights, bias terms, subsampling factors for pooling layers, number and size of kernels in different layers, number of feature maps, etc.) can be stored in data modules 1024 and/or 1032.

[0094] The results of the cuboid detection and keypoint location (e.g., the output of the cuboid detector 200) can be used by one or both of the processing modules 1024, 1028 for additional operations or processing. For example, the processing modules 1024, 1028 of the wearable display system 1000 can be programmed to perform additional applications described herein (such as applications in augmented reality, human-computer interaction (HCI), autonomous vehicles, drones, or robotics in general) based on the output of the cuboid detector 200.

Additional Aspects

[0095] In a 1st aspect, a system for cuboid detection and keypoint localization is disclosed. The system comprises: non-transitory memory configured to store: executable instructions, an image for cuboid detection, and a cuboid detector comprising: a plurality of convolutional layers and non-convolutional layers of a first convolutional neural network

(CNN) for generating a convolutional feature map from the image, a region proposal network (RPN) comprising a second CNN for determining, using the convolutional feature map, at least one region of interest (RoI) comprising a cuboid at a cuboid image location of the image, and a pooling layer and at least one regressor layer for determining, using the convolutional feature map and the RoI comprising the cuboid, a refined RoI at a refined cuboid image location and a representation of the cuboid; a hardware processor in communication with the non-transitory memory, the hardware processor programmed by the executable instructions to: receive the image; generate, using the plurality of convolutional layers and the non-convolutional layers of the first CNN and the image, the convolutional feature map; determine, using the RPN, the at least one RoI comprising the cuboid at the cuboid image location of the image; determine, using the pooling layer and the cuboid image location, a submap of the convolutional feature map corresponding to the RoI comprising the cuboid; and determine, using the at least one regressor layer and the submap of the convolutional feature map corresponding to the RoI comprising the cuboid, the refined RoI at the refined cuboid image location and the representation of the cuboid.

[0096] In a 2nd aspect, the system of aspect 1, wherein the hardware processor is further programmed to: determine, using the refined cuboid image location, a refined submap of the convolutional feature map corresponding to the refined RoI comprising the cuboid; determine, using the pooling layer, the at least one regressor layer, and the refined submap of the convolutional feature map corresponding to the refined RoI comprising the cuboid, a further refined RoI at a further refined cuboid image location and a further defined representation of the cuboid.

[0097] In a 3rd aspect, the system of any one of aspects 1-2, wherein the cuboid image location is represented as a two-dimensional (2D) bounding box.

[0098] In a 4th aspect, the system of any one of aspects 1-3, wherein the refined cuboid image location is represented as a two-dimensional (2D) bounding box.

[0099] In a 5th aspect, the system of any one of aspects 1-4, wherein the non-convolutional layers of the first CNN comprises a normalization layer, a brightness normalization layer, a batch normalization layer, a rectified linear layer, an upsampling layer, a concatenation layer, a pooling layer, a softsign layer, or any combination thereof.

[0100] In a 6th aspect, the system of any one of aspects 1-5, wherein the at least one regressor layer comprises two or more layers.

[0101] In a 7th aspect, the system of aspect 6, wherein the two or more layers comprise a fully connected layer, a non-fully connected layer, or any combination thereof.

[0102] In a 8th aspect, the system of any one of aspects 1-7, wherein the at least one regressor layer is associated with at least three loss functions during training.

[0103] In a 9th aspect, the system of aspect 8, wherein the at least three loss functions comprises a log loss function and a smooth L1 loss function.

[0104] In a 10th aspect, the system of any one of aspects 1-9, wherein RPN comprises a deep neural network (DNN).

[0105] In a 11th aspect, the system of any one of aspects 1-10, wherein the RPN is associated with at least two loss functions during training.

[0106] In a 12th aspect, the system of aspect 11, wherein the at least two loss functions comprises a log loss function and a smooth L1 loss function.

[0107] In a 13th aspect, the system of any one of aspects 1-12, wherein the representation of the cuboid comprises a parameterized representation of the cuboid.

[0108] In a 14th aspect, the system of aspect 13, wherein the parameterized representation of the cuboid comprises locations of a plurality of keypoints of the cuboid in the image.

[0109] In a 15th aspect, the system of aspect 14, wherein the plurality of keypoints comprises eight vertices of the cuboid in the image.

[0110] In a 16th aspect, the system of aspect 13, wherein the parameterized representation comprises normalized offsets of the plurality of keypoints of the cuboid from the center of the image.

[0111] In a 17th aspect, the system of aspect 13, wherein the parameterized representation comprises N tuples.

[0112] In a 18th aspect, the system of aspect 13, wherein the parameterized representation of the cuboid comprises 12 parameters.

[0113] In a 19th aspect, the system of aspect 13, wherein the parameterized representation of the cuboid comprises a vanishing point parameterization.

[0114] In a 20th aspect, the system of any one of aspects 1-19, wherein the hardware processor is further programmed to: interact with a user of the system based on the refined RoI at the refined cuboid image location and the representation of the cuboid.

[0115] In a 21st aspect, the system of aspect 20, wherein the cuboid corresponds to a stationary box, and wherein to interact with the user of the system, the hardware processor is further programmed to: generate character animation in relation to the stationary box based on the refined image location of the cuboid and the representation of the cuboid.

[0116] In a 22nd aspect, the system of aspect 20, wherein the cuboid corresponds to a hand-held cuboid, and wherein to interact with the user of the system, the hardware processor is further programmed to: determine a pose of the cuboid using the representation of the cuboid; and interact with the user of the system based on the pose of the cuboid.

[0117] In a 23rd aspect, the system of aspect 20, wherein the cuboid corresponds to a rare object not recognizable by a third CNN, and wherein to interact with the user of the system, the hardware processor is further programmed to: provide the user with a notification that the rare object not recognizable by the third CNN is detected.

[0118] In a 24th aspect, the system of any one of aspects 1-23, wherein the cuboid corresponds to a man-made structure, and wherein the hardware processor is further programmed to: assist a user of the system during an unmanned flight based on the refined RoI at the refined cuboid image location and the representation of the cuboid.

[0119] In a 25th aspect, the system of aspect any one of aspects 1-24, wherein the cuboid corresponds to a marker, and wherein the hardware processor is further programmed to: perform simultaneous location and mapping (SLAM) based on the refined RoI at the refined cuboid image location and the representation of the cuboid.

[0120] In a 26th aspect, a wearable display system is disclosed. The wearable display comprises: an outward-facing imaging system configured to obtain an image for cuboid detection; and the system for cuboid detection and keypoint localization of any one of aspects 1-25.

[0121] In a 27th aspect, a system for training a cuboid detector is disclosed. The system comprises: non-transitory memory configured to store executable instructions; and a hardware processor in communication with the non-transitory memory, the hardware processor programmed by the executable instructions to: receive a plurality of training

images each comprising at least one cuboid; generate a cuboid detector, wherein the cuboid detector comprises: a plurality of convolutional layers and non-convolutional layers of a first convolutional neural network (CNN), a region proposal network (RPN) connected to a first layer of the plurality of convolutional layers and non-convolutional layers, and a pooling layer and at least one regressor layer, the pooling layer and the at least one regressor layer connected to a second layer of the plurality of convolutional layers and non-convolutional layers; and train the cuboid detector, wherein to train the cuboid detector, the hardware process is configured to: determine, using the cuboid detector, a RoI at a cuboid image location and a representation of a cuboid in a training image of the plurality of training images; determine a first difference between a reference cuboid image location and the cuboid image location and a second difference between a reference representation of the cuboid and the determined representation of the cuboid; and update weights of the cuboid detector based on the first difference and the second difference.

[0122] In a 28th aspect, the system of aspect 27, wherein the cuboid comprises a cuboid, a cylinder, a sphere, or any combination thereof.

[0123] In a 29th aspect, the system of any one of aspects 27-28 wherein the first layer of the plurality of convolutional layers and non-convolutional layers and the second layer of the plurality of convolutional layers and non-convolutional layers are identical.

[0124] In a 30th aspect, the system of any one of aspects 27-29, wherein to determine the RoI at the cuboid image location and the representation of the cuboid, the hardware processor is further programmed to: generate, using the plurality of convolutional layers and the non-convolutional layers, a convolutional feature map for the at least one training image of the plurality of training images; determine, using the RPN, at least one RoI comprising the cuboid at an initial cuboid image location in the training image; determine, using the initial cuboid image location, a submap of the convolutional feature map corresponding to the at least one RoI comprising the cuboid; and determine, using the pooling layer, the at least one regressor layer, and the submap of the convolutional feature map corresponding to the at least one RoI comprising the cuboid, the RoI at the cuboid image location and the representation of the cuboid.

[0125] In a 31st aspect, the system of any one of aspects 27-30, wherein the initial cuboid image location is represented as a two-dimensional (2D) bounding box.

[0126] In a 32nd aspect, the system of any one of aspects 27-31, wherein to determine the RoI at the cuboid image location and the representation of the cuboid, the hardware processor is further programmed to: iteratively determine, using the pooling layer, the at least one regressor layer, and the submap of the convolutional feature map corresponding to the RoI comprising the cuboid, the RoI at the cuboid image location and the representation of the cuboid.

[0127] In a 33rd aspect, the system of any one of aspects 27-32, wherein the initial cuboid image location is represented as a two-dimensional (2D) bounding box.

[0128] In a 34th aspect, the system of any one of aspects 27-33, wherein to update weights of the cuboid detector, the hardware-based processor is programmed to: update the weights of the RPN and the weights of the at least one regressor layer.

[0129] In a 35th aspect, the system of any one of aspects 27-33, wherein to update weights of the cuboid detector, the hardware-based processor is programmed to: update the weights of the RPN and the weights of the at least one regressor layer without updating the weights of the first CNN.

[0130] In a 36th aspect, the system of any one of aspects 27-33, wherein to update weights of the cuboid detector, the hardware-based processor is programmed to: update the weights of the first CNN, the weights of the RPN, and the weights of the at least one regressor layer.

[0131] In a 37th aspect, the system of any one of aspects 27-36, wherein to generate the cuboid detector, the hardware-based processor is programmed to: receive the first CNN.

[0132] In a 38th aspect, the system of any one of aspects 27-37, wherein the at least one regressor layer comprises two or more layers.

[0133] In a 39th aspect, the system of aspect 38, wherein the two or more layers comprise a fully connected layer, a non-fully connected layer, or any combination thereof.

[0134] In a 40th aspect, the system of any one of aspects 27-38, wherein the at least one regressor layer is associated with at least three loss functions during training of the cuboid detector.

[0135] In a 41st aspect, the system of aspect 40, wherein the at least three loss functions comprises a log loss function and a smooth L1 loss function.

[0136] In a 42nd aspect, the system of any one of aspects 27-41, wherein RPN comprises a deep neural network (DNN).

[0137] In a 43rd aspect, the system of any one of aspects 27-42, wherein the RPN is associated with at least two loss functions during the training of the cuboid detector.

[0138] In a 44th aspect, the system of aspect 43, wherein the at least two loss functions comprises a log loss function and a smooth L1 loss function.

[0139] In a 45th aspect, the system of any one of aspects 27-44, wherein the representation of the cuboid comprises a parameterized representation of the cuboid.

[0140] In a 46th aspect, the system of aspect 45, wherein the parameterized representation comprises N tuples.

[0141] In a 47th aspect, a wearable display system is disclosed. The wearable display system comprises: an outward-facing imaging system configured to obtain an image of an environment of the wearer of the wearable display system; non-transitory memory configured to store the image; and a hardware processor in communication with the non-transitory memory, the processor programmed to: access the image of the environment; analyze the image to detect a cuboid in the image, wherein to analyze the image, the processor is programmed to: utilize layers of a convolutional neural network (CNN) to generate a convolutional feature map comprising features; utilize a region proposal network (RPN) to map the convolutional feature map into a region of interest (RoI); pool features in the RoI to generate first pooled features; pass the first pooled features through a regressor to generate a first bounding box estimate and a first cuboid vertex estimate; generate second pooled features based on the first bounding box estimate; and pass the second pooled features through the regressor to generate a second bounding box estimate and a second cuboid vertex estimate.

[0142] In a 48th aspect, the wearable display system of aspect 47, wherein the image comprises a monocular color image.

[0143] In a 49th aspect, the wearable display system of aspect 47 or aspect 48, wherein the RPN comprises a CNN that maps the convolutional feature map to the RoI.

[0144] In a 50th aspect, the wearable display system of any one of aspects 47 to 49, wherein the first bounding box estimate or the second bounding box estimate comprise offsets from a center of a bounding box.

[0145] In a 51st aspect, a system for detecting a cuboid in an image is disclosed. The system comprises: non-transitory memory configured to store an image of a region; a hardware processor in communication with the non-transitory memory, the processor programmed to: evaluate a convolutional neural network to generate a feature map; analyze the feature map to obtain a region of interest (RoI); determine that the RoI contains a cuboid; analyze first pooled features in the RoI of the feature map to generate a first estimate for vertices of the cuboid; generate an improved RoI based at least in part on the first estimate for the vertices of the cuboid; analyze second pooled features in the improved RoI of the feature map to generate a second estimate for vertices of the cuboid; and output the second estimate for vertices of the cuboid.

[0146] In a 52nd aspect, the system of aspect 51, wherein to analyze the feature map to obtain a region of interest (RoI), the processor is programmed to evaluate a region proposal network (RPN).

[0147] In a 53rd aspect, the system of aspect 51 or 52, wherein the first estimate for vertices of the cuboid comprise offsets from a center of the RoI, or the second estimate for vertices of the cuboid comprise offsets from a center of the improved RoI.

[0148] In a 54th aspect, a method for cuboid detection and keypoint localization is disclosed. The method is under control of a hardware processor and comprises: receiving an image; generating, using a plurality of convolutional layers and non-convolutional layers of a first convolutional neural network (CNN) of a cuboid detector and the image, a convolutional feature map; determining, using a region proposal network (RPN) comprising a second CNN of the cuboid detector, at least one RoI comprising a cuboid at a cuboid image location of the image; determining, using a pooling layer of the cuboid detector and the cuboid image location, a submap of the convolutional feature map corresponding to the RoI comprising the cuboid; and determining, using at least one regressor layer of the cuboid detector and the submap of the convolutional feature map corresponding to the RoI comprising the cuboid, a refined RoI at a refined cuboid image location and the representation of the cuboid.

[0149] In a 55th aspect, the method of aspect 54, further comprising: determining, using the refined cuboid image location, a refined submap of the convolutional feature map corresponding to the refined RoI comprising the cuboid; determining, using the pooling

layer, the at least one regressor layer, and the refined submap of the convolutional feature map corresponding to the refined RoI comprising the cuboid, a further refined RoI at a further refined cuboid image location and a further defined representation of the cuboid.

[0150] In a 56th aspect, the method of any one of aspects 54-55, wherein the cuboid image location is represented as a two-dimensional (2D) bounding box.

[0151] In a 57th aspect, the method of any one of aspects 54-56, wherein the refined cuboid image location is represented as a two-dimensional (2D) bounding box.

[0152] In a 58th aspect, the method of any one of aspects 54-57, wherein the non-convolutional layers of the first CNN comprises a normalization layer, a brightness normalization layer, a batch normalization layer, a rectified linear layer, an upsampling layer, a concatenation layer, a pooling layer, a softsign layer, or any combination thereof.

[0153] In a 59th aspect, the method of any one of aspects 54-58, wherein the at least one regressor layer comprises two or more layers.

[0154] In a 60th aspect, the method of aspect 59, wherein the two or more layers comprise a fully connected layer, a non-fully connected layer, or any combination thereof.

[0155] In a 61st aspect, the method of any one of aspects 54-60, wherein RPN comprises a deep neural network (DNN).

[0156] In a 62nd aspect, the method of any one of aspects 54-61, wherein the representation of the cuboid comprises a parameterized representation of the cuboid.

[0157] In a 63rd aspect, the method of aspect 62, wherein the parameterized representation of the cuboid comprises locations of a plurality of keypoints of the cuboid in the image.

[0158] In a 64th aspect, the method of aspect 63, wherein the plurality of keypoints comprises eight vertices of the cuboid in the image.

[0159] In a 65th aspect, the method of aspect 62, wherein the parameterized representation comprises normalized offsets of the plurality of keypoints of the cuboid from the center of the image.

[0160] In a 66th aspect, the method of aspect 62, wherein the parameterized representation comprises N tuples.

[0161] In a 67th aspect, the method of aspect 62, wherein the parameterized representation of the cuboid comprises 12 parameters.

[0162] In a 68th aspect, the method of aspect 62, wherein the parameterized representation of the cuboid comprises a vanishing point parameterization.

[0163] In a 69th aspect, the method of any one of aspects 54-58, further comprising: interacting with a user based on the refined RoI at the refined cuboid image location and the representation of the cuboid.

[0164] In a 70th aspect, the method of aspect 69, wherein the cuboid corresponds to a stationary box, and interacting with the user comprises: generating character animation in relation to the stationary box based on the refined image location of the cuboid and the representation of the cuboid.

[0165] In a 71st aspect, the method of aspect 69, wherein the cuboid corresponds to a hand-held cuboid, and wherein interacting with the user comprises: determining a pose of the cuboid using the representation of the cuboid; and interacting with the user based on the pose of the cuboid.

[0166] In a 72nd aspect, the method of aspect 69, wherein the cuboid corresponds to a rare object not recognizable by a third CNN, and wherein interacting with the user comprises: providing the user with a notification that the rare object not recognizable by the third CNN is detected.

[0167] In a 73rd aspect, the method of any one of aspects 54-72, further comprising: assisting a user of the system during an unmanned flight based on the refined RoI at the refined cuboid image location and the representation of the cuboid, wherein the cuboid corresponds to a man-made structure.

[0168] In a 74th aspect, the method of any one of aspects 54-73, further comprising: perform simultaneous location and mapping (SLAM) based on the refined RoI at the refined cuboid image location and the representation of the cuboid, wherein the cuboid corresponds to a marker.

[0169] In a 75th aspect, the method of any one of aspects 54-74, further comprising: receiving a plurality of training images each comprising at least one training cuboid; generating the cuboid detector and training the cuboid detector comprising: determining, using the cuboid detector, a training RoI at a training cuboid image location and a representation of a training cuboid in a training image of the plurality of training images; determining a first difference between a reference cuboid image location and the training

cuboid image location and a second difference between a reference representation of the training cuboid and the determined representation of the training cuboid; and updating weights of the cuboid detector based on the first difference and the second difference.

[0170] In a 76th aspect, the method of aspect 75, wherein determining the training RoI at the training cuboid image location and the representation of the training cuboid comprises: generating, using the plurality of convolutional layers and the non-convolutional layers, a training convolutional feature map for the at least one training image of the plurality of training images; determining, using the RPN, at least one training RoI comprising the training cuboid at an initial training cuboid image location in the training image; determining, using the initial training cuboid image location, a submap of the convolutional feature map corresponding to the at least one RoI comprising the cuboid; and determining, using the pooling layer, the at least one regressor layer, and the submap of the training convolutional feature map corresponding to the at least one training RoI comprising the training cuboid, the training RoI at the training cuboid image location and the representation of the training cuboid.

[0171] In a 77th aspect, the method of aspect 76, wherein the initial training cuboid image location is represented as a two-dimensional (2D) bounding box.

[0172] In a 78th aspect, the method of aspect 75, wherein determining the training RoI at the training cuboid image location and the representation of the training cuboid comprises: iteratively determining, using the pooling layer, the at least one regressor layer, and the submap of the training convolutional feature map corresponding to the training RoI comprising the training cuboid, the RoI at the training cuboid image location and the representation of the training cuboid.

[0173] In a 79th aspect, the method of aspect 78, wherein the initial training cuboid image location is represented as a two-dimensional (2D) bounding box.

[0174] In a 80th aspect, the method of any one of aspects 75-79, wherein updating weights of the cuboid detector comprises: updating the weights of the RPN and the weights of the at least one regressor layer.

[0175] In a 81st aspect, the method of any one of aspects 75-79, wherein updating weights of the cuboid detector comprises: updating the weights of the RPN and the weights of the at least one regressor layer without updating the weights of the first CNN.

[0176] In a 82nd aspect, the method of any one of aspects 75-79, wherein updating weights of the cuboid detector comprises: updating the weights of the first CNN, the weights of the RPN, and the weights of the at least one regressor layer.

[0177] In a 83rd aspect, the method of any one of aspects 54-82, wherein generating the cuboid detector comprises: receiving the first CNN.

[0178] In a 84th aspect, the method of any one of aspects 75-83, wherein the at least one regressor layer is associated with at least three loss functions during training of the cuboid detector.

[0179] In a 85th aspect, the method of claim 84, wherein the at least three loss functions comprises a log loss function and a smooth L1 loss function.

[0180] In a 86th aspect, the method of any one of aspects 75-85, wherein the RPN is associated with at least two loss functions during the training of the cuboid detector.

[0181] In a 87th aspect, the method of claim 86, wherein the at least two loss functions comprises a log loss function and a smooth L1 loss function.

[0182] In a 88th aspect, a method is disclosed. The method is under control of a hardware processor and comprises: accessing an image of the environment; analyzing the image to detect a cuboid in the image comprising utilizing layers of a convolutional neural network (CNN) to generate a convolutional feature map comprising features; utilizing a region proposal network (RPN) to map the convolutional feature map into a region of interest (RoI); pooling features in the RoI to generate first pooled features; passing the first pooled features through a regressor to generate a first bounding box estimate and a first cuboid vertex estimate; generating second pooled features based on the first bounding box estimate; and passing the second pooled features through the regressor to generate a second bounding box estimate and a second cuboid vertex estimate.

[0183] In a 89th aspect, the method of aspect 88, wherein the image comprises a monocular color image.

[0184] In a 90th aspect, the method of aspect 88 or aspect 89, wherein the RPN comprises a CNN that maps the convolutional feature map to the RoI.

[0185] In a 91st aspect, the method of any one of aspects 88 to 89, wherein the first bounding box estimate or the second bounding box estimate comprise offsets from a center of a bounding box.

[0186] In a 92nd aspect, a method for detecting a cuboid in an image is disclosed. The method is under control of a hardware processor and comprises: evaluating a convolutional neural network to generate a feature map; analyzing the feature map to obtain a region of interest (RoI); determining that the RoI contains a cuboid; analyzing first pooled features in the RoI of the feature map to generate a first estimate for vertices of the cuboid; generating an improved RoI based at least in part on the first estimate for the vertices of the cuboid; analyzing second pooled features in the improved RoI of the feature map to generate a second estimate for vertices of the cuboid; and outputting the second estimate for vertices of the cuboid.

[0187] In a 93rd aspect, the method of aspect 92, wherein analyzing the feature map to obtain a region of interest (RoI) comprises evaluating a region proposal network (RPN).

[0188] In a 94th aspect, the method of aspect 92 or 93, wherein the first estimate for vertices of the cuboid comprise offsets from a center of the RoI, or the second estimate for vertices of the cuboid comprise offsets from a center of the improved RoI.

Conclusion

[0189] Each of the processes, methods, and algorithms described herein and/or depicted in the attached figures may be embodied in, and fully or partially automated by, code modules executed by one or more physical computing systems, hardware computer processors, application-specific circuitry, and/or electronic hardware configured to execute specific and particular computer instructions. For example, computing systems can include general purpose computers (e.g., servers) programmed with specific computer instructions or special purpose computers, special purpose circuitry, and so forth. A code module may be compiled and linked into an executable program, installed in a dynamic link library, or may be written in an interpreted programming language. In some implementations, particular operations and methods may be performed by circuitry that is specific to a given function.

[0190] Further, certain implementations of the functionality of the present disclosure are sufficiently mathematically, computationally, or technically complex that application-specific hardware or one or more physical computing devices (utilizing appropriate specialized executable instructions) may be necessary to perform the

functionality, for example, due to the volume or complexity of the calculations involved or to provide results substantially in real-time. For example, a video may include many frames, with each frame having millions of pixels, and specifically programmed computer hardware is necessary to process the video data to provide a desired image processing task or application in a commercially reasonable amount of time.

[0191] Code modules or any type of data may be stored on any type of non-transitory computer-readable medium, such as physical computer storage including hard drives, solid state memory, random access memory (RAM), read only memory (ROM), optical disc, volatile or non-volatile storage, combinations of the same and/or the like. The methods and modules (or data) may also be transmitted as generated data signals (e.g., as part of a carrier wave or other analog or digital propagated signal) on a variety of computer-readable transmission mediums, including wireless-based and wired/cable-based mediums, and may take a variety of forms (e.g., as part of a single or multiplexed analog signal, or as multiple discrete digital packets or frames). The results of the disclosed processes or process steps may be stored, persistently or otherwise, in any type of non-transitory, tangible computer storage or may be communicated via a computer-readable transmission medium.

[0192] Any processes, blocks, states, steps, or functionalities in flow diagrams described herein and/or depicted in the attached figures should be understood as potentially representing code modules, segments, or portions of code which include one or more executable instructions for implementing specific functions (e.g., logical or arithmetical) or steps in the process. The various processes, blocks, states, steps, or functionalities can be combined, rearranged, added to, deleted from, modified, or otherwise changed from the illustrative examples provided herein. In some embodiments, additional or different computing systems or code modules may perform some or all of the functionalities described herein. The methods and processes described herein are also not limited to any particular sequence, and the blocks, steps, or states relating thereto can be performed in other sequences that are appropriate, for example, in serial, in parallel, or in some other manner. Tasks or events may be added to or removed from the disclosed example embodiments. Moreover, the separation of various system components in the implementations described herein is for illustrative purposes and should not be understood as requiring such separation in all implementations. It should be understood that the described program components, methods,

and systems can generally be integrated together in a single computer product or packaged into multiple computer products. Many implementation variations are possible.

[0193] The processes, methods, and systems may be implemented in a network (or distributed) computing environment. Network environments include enterprise-wide computer networks, intranets, local area networks (LAN), wide area networks (WAN), personal area networks (PAN), cloud computing networks, crowd-sourced computing networks, the Internet, and the World Wide Web. The network may be a wired or a wireless network or any other type of communication network.

[0194] The systems and methods of the disclosure each have several innovative aspects, no single one of which is solely responsible or required for the desirable attributes disclosed herein. The various features and processes described herein may be used independently of one another, or may be combined in various ways. All possible combinations and subcombinations are intended to fall within the scope of this disclosure. Various modifications to the implementations described in this disclosure may be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other implementations without departing from the spirit or scope of this disclosure. Thus, the claims are not intended to be limited to the implementations shown herein, but are to be accorded the widest scope consistent with this disclosure, the principles and the novel features disclosed herein.

[0195] Certain features that are described in this specification in the context of separate implementations also can be implemented in combination in a single implementation. Conversely, various features that are described in the context of a single implementation also can be implemented in multiple implementations separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination. No single feature or group of features is necessary or indispensable to each and every embodiment.

[0196] Conditional language used herein, such as, among others, “can,” “could,” “might,” “may,” “e.g.,” and the like, unless specifically stated otherwise, or otherwise

understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements and/or steps. Thus, such conditional language is not generally intended to imply that features, elements and/or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without author input or prompting, whether these features, elements and/or steps are included or are to be performed in any particular embodiment. The terms “comprising,” “including,” “having,” and the like are synonymous and are used inclusively, in an open-ended fashion, and do not exclude additional elements, features, acts, operations, and so forth. Also, the term “or” is used in its inclusive sense (and not in its exclusive sense) so that when used, for example, to connect a list of elements, the term “or” means one, some, or all of the elements in the list. In addition, the articles “a,” “an,” and “the” as used in this application and the appended claims are to be construed to mean “one or more” or “at least one” unless specified otherwise.

[0197] As used herein, a phrase referring to “at least one of” a list of items refers to any combination of those items, including single members. As an example, “at least one of: A, B, or C” is intended to cover: A, B, C, A and B, A and C, B and C, and A, B, and C. Conjunctive language such as the phrase “at least one of X, Y and Z,” unless specifically stated otherwise, is otherwise understood with the context as used in general to convey that an item, term, etc. may be at least one of X, Y or Z. Thus, such conjunctive language is not generally intended to imply that certain embodiments require at least one of X, at least one of Y and at least one of Z to each be present.

[0198] Similarly, while operations may be depicted in the drawings in a particular order, it is to be recognized that such operations need not be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. Further, the drawings may schematically depict one more example processes in the form of a flowchart. However, other operations that are not depicted can be incorporated in the example methods and processes that are schematically illustrated. For example, one or more additional operations can be performed before, after, simultaneously, or between any of the illustrated operations. Additionally, the operations may be rearranged or reordered in other implementations. In certain circumstances, multitasking and parallel

processing may be advantageous. Moreover, the separation of various system components in the implementations described above should not be understood as requiring such separation in all implementations, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products. Additionally, other implementations are within the scope of the following claims. In some cases, the actions recited in the claims can be performed in a different order and still achieve desirable results.

WHAT IS CLAIMED IS:

1. A system for cuboid detection and keypoint localization comprising:
 - non-transitory memory configured to store:
 - executable instructions,
 - an image for cuboid detection, and
 - a cuboid detector comprising:
 - a plurality of convolutional layers and non-convolutional layers of a first convolutional neural network (CNN) for generating a convolutional feature map from the image,
 - a region proposal network (RPN) comprising a second CNN for determining, using the convolutional feature map, at least one region of interest (RoI) comprising a cuboid at a cuboid image location of the image, and
 - a pooling layer and at least one regressor layer for determining, using the convolutional feature map and the RoI comprising the cuboid, a refined RoI at a refined cuboid image location and a representation of the cuboid;
 - a hardware processor in communication with the non-transitory memory, the hardware processor programmed by the executable instructions to:
 - receive the image;
 - generate, using the plurality of convolutional layers and the non-convolutional layers of the first CNN and the image, the convolutional feature map;
 - determine, using the RPN, the at least one RoI comprising the cuboid at the cuboid image location of the image;
 - determine, using the pooling layer and the cuboid image location, a submap of the convolutional feature map corresponding to the RoI comprising the cuboid; and
 - determine, using the at least one regressor layer and the submap of the convolutional feature map corresponding to the RoI comprising the cuboid, the refined RoI at the refined cuboid image location and the representation of

the cuboid.

2. The system of claim 1, wherein the hardware processor is further programmed to:

determine, using the refined cuboid image location, a refined submap of the convolutional feature map corresponding to the refined RoI comprising the cuboid;

determine, using the pooling layer, the at least one regressor layer, and the refined submap of the convolutional feature map corresponding to the refined RoI comprising the cuboid, a further refined RoI at a further refined cuboid image location and a further defined representation of the cuboid.

3. The system of claim 1, wherein the cuboid image location is represented as a two-dimensional (2D) bounding box.

4. The system of claim 1, wherein the refined cuboid image location is represented as a two-dimensional (2D) bounding box.

5. The system of claim 1, wherein the at least one regressor layer comprises two or more layers.

6. The system of claim 5, wherein the two or more layers comprise a fully connected layer, a non-fully connected layer, or any combination thereof.

7. The system of claim 1, wherein RPN comprises a deep neural network (DNN).

8. The system of claim 1, wherein the representation of the cuboid comprises a parameterized representation of the cuboid.

9. The system of claim 8, wherein the parameterized representation of the cuboid comprises locations of a plurality of keypoints of the cuboid in the image.

10. The system of claim 9, wherein the plurality of keypoints comprises eight vertices of the cuboid in the image.

11. The system of claim 8, wherein the parameterized representation comprises normalized offsets of the plurality of keypoints of the cuboid from the center of the image.

12. The system of claim 8, wherein the parameterized representation of the cuboid comprises 12 parameters.

13. The system of claim 8, wherein the parameterized representation of the cuboid comprises a vanishing point parameterization.

14. The system of claim 1, wherein the hardware processor is further programmed to:

interact with a user of the system based on the refined RoI at the refined cuboid image location and the representation of the cuboid.

15. The system of claim 14,
wherein the cuboid corresponds to a stationary box, and
wherein to interact with the user of the system, the hardware processor is further programmed to:

generate character animation in relation to the stationary box based on the refined image location of the cuboid and the representation of the cuboid.

16. The system of claim 14,
wherein the cuboid corresponds to a hand-held cuboid, and
wherein to interact with the user of the system, the hardware processor is further programmed to:

determine a pose of the cuboid using the representation of the cuboid;
and

interact with the user of the system based on the pose of the cuboid.

17. The system of claim 14,
wherein the cuboid corresponds to a rare object not recognizable by a third CNN, and

wherein to interact with the user of the system, the hardware processor is further programmed to:

provide the user with a notification that the rare object not recognizable by the third CNN is detected.

18. The system of claim 1,
wherein the cuboid corresponds to a man-made structure, and
wherein the hardware processor is further programmed to:

assist a user of the system during an unmanned flight based on the refined RoI at the refined cuboid image location and the representation of the cuboid.

19. The system of claim 1,

wherein the cuboid corresponds to a marker, and

wherein the hardware processor is further programmed to:

perform simultaneous location and mapping (SLAM) based on the refined RoI at the refined cuboid image location and the representation of the cuboid.

20. A method for cuboid detection and keypoint localization comprising, under control of a hardware processor:

receiving an image;

generating, using a plurality of convolutional layers and non-convolutional layers of a first convolutional neural network (CNN) of a cuboid detector and the image, a convolutional feature map;

determining, using a region proposal network (RPN) comprising a second CNN of the cuboid detector, at least one RoI comprising a cuboid at a cuboid image location of the image;

determining, using a pooling layer of the cuboid detector and the cuboid image location, a submap of the convolutional feature map corresponding to the RoI comprising the cuboid; and

determining, using at least one regressor layer of the cuboid detector and the submap of the convolutional feature map corresponding to the RoI comprising the cuboid, a refined RoI at a refined cuboid image location and the representation of the cuboid.

21. The method of claim 20, further comprising:

determining, using the refined cuboid image location, a refined submap of the convolutional feature map corresponding to the refined RoI comprising the cuboid;

determining, using the pooling layer, the at least one regressor layer, and the refined submap of the convolutional feature map corresponding to the refined RoI comprising the cuboid, a further refined RoI at a further refined cuboid image location and a further defined representation of the cuboid.

22. The method of claim 20, wherein the cuboid image location is represented as a two-dimensional (2D) bounding box, and wherein the refined cuboid image location is represented as a two-dimensional (2D) bounding box.

23. The method of claim 20, wherein the non-convolutional layers of the first CNN comprises a normalization layer, a brightness normalization layer, a batch normalization layer, a rectified linear layer, an upsampling layer, a concatenation layer, a pooling layer, a softsign layer, or any combination thereof.

24. The method of claim 20, wherein the at least one regressor layer comprises two or more layers, and wherein the two or more layers comprise a fully connected layer, a non-fully connected layer, or any combination thereof.

25. The method of claim 20, wherein the representation of the cuboid comprises a parameterized representation of the cuboid comprising locations of a plurality of keypoints of the cuboid in the image.




FIG. 1B

FIG. 1A




FIG. 2




FIG. 3




FIG. 4A




FIG. 4B




FIG. 4C




FIG. 4D




FIG. 4E




FIG. 4F




FIG. 4G

Iteration 1

Iteration 2




FIG. 5A

Iteration 1

Iteration 2




FIG. 5B

Iteration 1

Iteration 2




FIG. 5C






FIG. 6




APK vs. Recall

FIG. 7B




Precision-Recall Curve

FIG. 7A



Keypoint-wise APK

FIG. 7D



PCK vs. α

FIG. 7C




FIG. 7F

FIG. 7E




FIG. 8




FIG. 9




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2017/061618

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 3/01; G06F 3/0346; G06K 9/66; G06N 3/04; G06T 7/215 (2018.01)

CPC - G06F 1/163; G06F 3/0346; G06K 9/00979; G06K 9/4628; G06K 9/4671; G06N 3/0454; G06T 19/006; H04N 5/145 (2018.01)

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)

See Search History document

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

USPC - 382/103; 382/155; 382/156; 382/190; 382/279 (keyword delimited)

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

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	CHEN et al. "3d object proposals for accurate object class detection." In: Advances in Neural Information Processing Systems. 2015 (2015) Retrieved from <http://papers.nips.cc/paper/5644-3d-object-proposals-for-accurate-object-class-detection.pdf> entire document	1-25
A	US 2016/0104056 A1 (MICROSOFT TECHNOLOGY LICENSING, LLC) 14 April 2016 (14.04.2016) entire document	1-25
A	DAI et al. "R-fcn: Object detection via region-based fully convolutional networks." In: Advances in neural information processing systems. 21 June 2016 (21.06.2016) Retrieved from <https://arxiv.org/pdf/1605.06409.pdf> entire document	1-25
A	GUPTA et al. "Learning rich features from RGB-D images for object detection and segmentation." In: European Conference on Computer Vision. 22 July 2014 (22.07.2014) Retrieved from <https://arxiv.org/pdf/1407.5736.pdf> entire document	1-25
P,A	PAVLAKOS et al. "6-dof object pose from semantic keypoints." In: arXiv preprint. 14.03.2017. retrieved on 03.01.2018. Retrieved from <http://www.cis.upenn.edu/~kostas/mypub.dir/pavlakos17icra.pdf> entire document	1-25



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

31 December 2017

Date of mailing of the international search report

17 JAN 2018

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, VA 22313-1450
Facsimile No. 571-273-8300

Authorized officer

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774