

(51) International Patent Classification:
G06K 9/62 (2006.01)

(21) International Application Number:

PCT/US2010/040678

(22) International Filing Date:

30 June 2010 (30.06.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

61/222,245 1 July 2009 (01.07.2009) US

(71) Applicant (for all designated States except US): **HONDA MOTOR CO, LTD.** [JP/JP]; 2-1-1 Minami Aoyama, Minato-ku, Tokyo, 107-8556 (JP).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **HEISELE, Bernd** [DE/US]; c/o Honda Patents & Technologies North America LLC, 21001 State Route 739, Raymond, OH 43067-9705 (US). **KIM, Gunhee** [KR/US]; c/o Honda Patents & Technologies North America LLC, 21001 State Route 739, Raymond, OH 43067-9705 (US). **MEYER,****Andrew, J.** [US/US]; C/o Honda Patents & Technologies North America LLC, 21001 State Route 739, Raymond, OH 43067-9705 (US).(74) Agent: **MCNELIS, John;** Fenwick & West LLP, 801 California Street, Mountain View, CA 94041 (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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, 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, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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,

[Continued on next page]

(54) Title: OBJECT RECOGNITION WITH 3D MODELS




FIG. 4

(57) Abstract: An "active learning" method trains a compact classifier for view-based object recognition. The method actively generates its own training data. Specifically, the generation of synthetic training images is controlled within an iterative training process. Valuable and/or informative object views are found in a low-dimensional rendering space and then added iteratively to the training set. In each iteration, new views are generated. A sparse training set is iteratively generated by searching for local minima of a classifier's output in a low-dimensional space of rendering parameters. An initial training set is generated. The classifier is trained using the training set. Local minima are found of the classifier's output in the low-dimensional rendering space. Images are rendered at the local minima. The newly-rendered images are added to the training set. The procedure is repeated so that the classifier is retrained using the modified training set.

WO 2011/002938 A1



LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, **Published:**

SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, — *with international search report (Art. 21(3))*
GW, ML, MR, NE, SN, TD, TG).

OBJECT RECOGNITION WITH 3D MODELS

INVENTORS

BERND HEISELE
GUNHEE KIM
ANDREW J. MEYER

RELATED APPLICATION

[0001] This application claims priority from U.S. provisional application number 61/222,245, filed July 1, 2009, entitled “Object Recognition with 3D Models”, which is incorporated by reference herein in its entirety.

FIELD OF THE INVENTION

[0002] The present invention relates to machine vision, including object recognition. More particularly, the present invention relates to training classifiers to perform view-based object recognition.

BACKGROUND OF THE INVENTION

[0003] Over the past few years, the object recognition community has taken on the challenge of developing systems that learn to recognize hundreds of object classes from a few examples per class. The standard data sets used for benchmarking these systems contained, in average, fewer than two hundred images per class (as opposed to sets of thousands of meticulously segmented object images that were used in earlier work on object detection). Benchmarking on such small data sets is inherently problematic. The test results cannot be generalized and can be misleading.

[0004] There have been efforts of building larger databases of manually annotated, natural images. However, the somewhat arbitrary selection of images and the missing ground truth make it difficult to systematically analyze specific properties of object recognition systems, such as invariance to pose, scale, position, and illumination. A database for shape-based object recognition which addresses these issues is NORB from the Courant Institute at New York University. Pictures of objects were taken with consideration of viewpoint and illumination. The images were synthetically altered to add more image variations, such as object rotation, background, and “distractors”.

[0005] Taking the idea of controlling the image generation one step further takes us to fully synthetic images rendered from realistic three-dimensional (3D) computer graphics models. Some view-based object recognition systems have been trained and evaluated on synthetic images. At least one face recognition system and one object recognition system have been trained on views of 3D models and tested on real images. 3D models have also been used in a generative approach to object recognition in which rendering parameter values are optimized such that the synthetic image best matches a given photographic image. To avoid getting trapped in local minima, this analysis-through-synthesis approach requires a good initial estimate of the rendering parameter values, making it unsuited to many object recognition/detection tasks.

SUMMARY OF THE INVENTION

[0006] A view-based object recognition system uses a model of an object's appearance in order to determine whether the object is present in a given image. The view-based object recognition system generates ("learns") the model when the system is trained using a training data set. The training data set includes images that are positive examples (i.e., the target object is present) and images that are negative examples (i.e., the target object is absent). Each training image is labeled correctly regarding whether it is a positive example or a negative example.

[0007] Like any other image, a training image can be either natural or synthetic. A natural image faithfully represents the appearance of a real-world object and is generated by, for example, taking a picture of the object using a camera or other image sensor. A synthetic image is any image other than a natural image. For example, a computer-aided design (CAD) model is rendered to generate the synthetic image. Training a view-based object recognition system using synthetic images (as opposed to natural images) has several advantages.

[0008] An "active learning" method is presented, in which the generation of synthetic training images is controlled within an iterative training process. (The term "active learning" is used herein to refer to a learning method that actively generates its own training data.) The primary idea of active learning is to find valuable and/or informative object views in a low-dimensional rendering space and then add these views iteratively to the training set. In each iteration, instead of "bootstrapping" the classifier (by adding particular sample images from a given database of images), new views are generated. A sparse training set is iteratively

generated by searching for local minima of a classifier's output in a low-dimensional space of rendering parameters.

[0009] In one embodiment, an active learning method trains a compact classifier for view-based object recognition. An initial training set is generated. The classifier is trained using the training set. Local minima are found of the classifier's output in the low-dimensional rendering space. Images are rendered at the local minima. The newly-rendered images are added to the training set. The procedure is repeated so that the classifier is retrained using the modified training set.

[0010] The features and advantages described in the specification are not all inclusive and, in particular, many additional features and advantages will be apparent to one of ordinary skill in the art in view of the drawings, specification, and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] FIG. 1 is a flowchart of an active learning method that trains a compact classifier for view-based object recognition, according to one embodiment of the invention.

[0012] FIG. 2 shows five computer graphics models, according to one embodiment of the invention.

[0013] FIG. 3 shows photographs of the printed versions of the five models of FIG. 2, according to one embodiment of the invention.

[0014] FIG. 4 is a block diagram of a system for performing the active learning method of FIG. 1, according to one embodiment of the invention.

[0015] FIG. 5 shows examples of rendered images after scaling and smoothing, according to one embodiment of the invention.

[0016] FIG. 6 shows examples of initially-selected views and views at nearby local minima, according to one embodiment of the invention.

[0017] The figures depict various embodiments of the present invention for purposes of illustration only. One skilled in the art will readily recognize from the following discussion that alternative embodiments of the structures and methods illustrated herein may be employed without departing from the principles of the invention described herein.

DETAILED DESCRIPTION OF THE INVENTION

[0018] Various embodiments are now described with reference to the figures where like reference numbers indicate identical or functionally similar elements. Also in the figures, the left-most digit of each reference number corresponds to the figure in which the reference number is first used.

[0019] Reference in the specification to “one embodiment” or to “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiments is included in at least one embodiment of the invention. The appearances of the phrase “in one embodiment” or “an embodiment” in various places in the specification are not necessarily all referring to the same embodiment.

[0020] Some portions of the detailed description that follows are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps (instructions) leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical, magnetic, or optical signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It is convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like. Furthermore, it is also convenient at times to refer to certain arrangements of steps requiring physical manipulations or transformation of physical quantities or representations of physical quantities as modules or code devices, without loss of generality.

[0021] However, all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as “processing” or “computing” or “calculating” or “determining” or “displaying” or the like refer to the action and processes of a computer system, or similar electronic computing device (such as a specific computing machine), that manipulates and transforms data represented as physical (electronic) quantities

within the computer system memories or registers or other such information storage, transmission, or display devices.

[0022] Certain aspects of the present invention include process steps and instructions described herein in the form of an algorithm. It should be noted that the process steps and instructions of the present invention could be embodied in software, firmware, or hardware and, when embodied in software, could be downloaded to reside on and be operated from different platforms used by a variety of operating systems. The invention can also be in a computer program product that can be executed on a computing system.

[0023] The present invention also relates to an apparatus for performing the operations herein. This apparatus may be specially constructed for the purposes, e.g., a specific computer, or it may comprise a general-purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium such as, but not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, application specific integrated circuits (ASICs), or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus. Memory can include any of the above and/or other devices that can store information/data/programs. Furthermore, the computers referred to in the specification may include a single processor or may be architectures employing multiple processor designs for increased computing capability.

[0024] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general-purpose systems may also be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the method steps. The structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the present invention as described herein, and any references below to specific languages are provided for disclosure of enablement and best mode of the present invention.

[0025] In addition, the language used in the specification has been principally selected for readability and instructional purposes and may not have been selected to delineate or circumscribe the inventive subject matter. Accordingly, the disclosure of the present

invention is intended to be illustrative, but not limiting, of the scope of the invention, which is set forth in the claims.

[0026] While particular embodiments and applications of the present invention have been illustrated and described herein, it is to be understood that the invention is not limited to the precise construction and components disclosed herein and that various modifications, changes, and variations may be made in the arrangement, operation, and details of the methods and apparatuses of the present invention without departing from the spirit and scope of the invention as it is defined in the appended claims.

[0027] A view-based object recognition system uses a model of an object's appearance in order to determine whether the object is present in a given image. The model represents the object's appearance in an image under various conditions (e.g., different poses and/or illuminations). The view-based object recognition system generates ("learns") the model when the system is trained using a training data set. The training data set includes images that are positive examples (i.e., the target object is present) and images that are negative examples (i.e., the target object is absent). Each training image is labeled correctly regarding whether it is a positive example or a negative example.

[0028] Like any other image, a training image can be either natural or synthetic. A natural image faithfully represents the appearance of a real-world object and is generated by, for example, taking a picture of the object using a camera or other image sensor. A synthetic image is any image other than a natural image. One type of synthetic image is generated based on a natural image. For example, the natural image is distorted or transformed in some way to generate the synthetic image. Another type of synthetic image is generated independently of a natural image. For example, a computer-aided design (CAD) model is rendered to generate the synthetic image. In order to distinguish between these two types of synthetic images, the first type will be referred to as partially-synthetic (since it was generated based on a natural image), and the second type will be referred to as fully-synthetic (since it was generated independently of a natural image).

[0029] Training a view-based object recognition system using synthetic images (as opposed to natural images) has several advantages. First of all, large numbers of training images can be generated automatically. Also, full control is available over image generation parameters, including internal and external camera parameters, illumination, composition of the scene, and animation of the scene. Finally, ground truth is known for the location, scale, and

orientation of each object. In video sequences, the frame rate, the camera motion, and the motion of objects are known.

[0030] Supervised machine learning methods, such as statistical classifiers, have been used to address the problem of view-based object recognition. Consider training a view-based classifier using fully-synthetic training images (e.g., images generated based on three-dimensional (3D) models). The training will likely face two problems. First, the training set will be very large, which can break the learning algorithm. Second, the solutions are not sparse, so the classifiers will be slow at run-time. One of the largest groups of modern machine learning algorithms is the group of kernel-based machine learning algorithms, which includes support vector machines (amongst others). It is possible to compute a quadratic matrix of size $N \times N$ where N is the number of samples that can “break” a kernel-based machine learning algorithm (see O. Chapelle, “Training a Support Vector Machine in the Primal”, Neural Computation, Vol. 19, No. 5, pp. 1155-1178, March 2007). A large number of samples will also make training slow, since the time complexity in kernel algorithms is polynomial in the number of samples. The complexity at runtime is also affected for support vector machines (and possibly for other kernel-based algorithms as well). The number of training data points that are used in computing the classification result increases linearly with the size of the training set. Basically, runtime complexity during classification is linear in training sample size. What is needed is a technique that can train (e.g., build or compute) compact classifiers from sparse data sets.

[0031] An “active learning” method is now presented, in which the generation of synthetic training images is controlled within an iterative training process. (Note that the term “active learning” in the machine learning community refers to learning from data with hidden labels. The term “active learning” is used herein to refer to something different – namely, a learning method that actively generates its own training data.) The primary idea of active learning is to find valuable and/or informative object views in a low-dimensional rendering space and then add these views iteratively to the training set. In each iteration, instead of “bootstrapping” the classifier (by adding particular sample images from a given database of images), new views are generated. A sparse training set is iteratively generated by searching for local minima of a classifier’s output in a low-dimensional space of rendering parameters.

[0032] FIG. 1 is a flowchart of an active learning method that trains a compact classifier for view-based object recognition, according to one embodiment of the invention. In step 110, an initial training set is generated. In step 120, the classifier is trained using the training set. In

step 130, local minima are found of the classifier's output in the low-dimensional rendering space. In step 140, images are rendered at the local minima. In step 150, the newly-rendered images are added to the training set. In step 160, the procedure is repeated starting from step 120, so that the classifier is retrained using the modified training set.

[0033] Steps 110-160 are further explained below in conjunction with an example of a classification task and a system for performing the active learning method of FIG. 1. In this example, object recognition is performed for five classes of objects. Each object is represented by a texture-less 3D model with 30,000 surface triangles. FIG. 2 shows the five computer graphics models, according to one embodiment of the invention. In FIG. 2, the illustrated models represent (from left to right) a bust of the goddess Aphrodite, a bear lying down, an elephant standing up, a horse standing up, and a shoe sole. Each of the five models is also printed on a 3D printer in order to be able to test the object recognition system on natural images. FIG. 3 shows photographs of the five printed models, according to one embodiment of the invention.

[0034] In this example, to generate an image (either synthetic or natural), a camera viewpoint (either virtual or actual) is moved on a sphere around the model, pointing towards the model's center. The model is illuminated by a point light source (which can be positioned on a sphere around the model) and by ambient light. This image generation setup results in a low-dimensional rendering space (specifically, a six-dimensional rendering space). The six rendering parameters are: 1) the viewpoint's location in azimuth, 2) the viewpoint's location in elevation, 3) the viewpoint's rotation around its optical axis, 4) the point light source's location in azimuth, 5) the point light source's location in elevation, and 6) the intensity ratio between the ambient light and the point light source.

[0035] FIG. 4 is a block diagram of a system for performing the active learning method of FIG. 1, according to one embodiment of the invention. The system 400 is able to train a compact classifier for view-based object recognition. The illustrated system 400 includes an active learning module 405, a renderer 410, a classifier trainer 415, a classifier tester 420, storage 425, and operating system tools 430.

[0036] In one embodiment, the active learning module 405 (and its component modules), the renderer 410, the classifier trainer 415, the classifier tester 420, and the operating system tools 430 are one or more computer program modules stored on one or more computer readable storage mediums and executing on one or more processors. The storage 425 (and its

contents) is stored on one or more computer readable storage mediums. Additionally, the active learning module 405 (and its component modules), the renderer 410, the classifier trainer 415, the classifier tester 420, the operating system tools 430, and the storage 425 are communicatively coupled to one another to at least the extent that data can be passed between them. In one embodiment, the operating system tools 430 are executing as one or more parts of an operating system on a personal computer, and the active learning module 405, the renderer 410, the classifier trainer 415, and the classifier tester 420 are executing on the same personal computer.

[0037] The illustrated storage 425 stores a classifier 445, a training set 450, and one or more 3D models 460. The classifier 445 is a statistical classifier such as a support vector machine (SVM) or a nearest-neighbor classifier. In one embodiment, the classifier 445 is a SVM with a Gaussian kernel. The SVM parameter values are optimized in initial experiments and then kept fixed throughout (e.g., $s = 2.0$ and $C = 10$, where s represents the kernel width and C represents the error cost). SVMs are further described in V. Vapnik, *Statistical Learning Theory*, Wiley-Interscience, 1998. Implementations of SVMs include, for example, the LIBSVM library (C.-C. Chang et al., “LIBSVM: a Library for Support Vector Machines”, 2001).

[0038] The classifier 445 operates using a vector of image feature values (“feature vector”). Any image feature can be used, such as gray-values, normalized gray-values, or histograms of gradients. In one embodiment, orientation histograms or histograms of gradients are used for the feature vector. The histograms are computed at five fixed locations within a 23x23 image, resulting in a 640-dimensional feature vector that is normalized to unit length. The 128-dimensional histograms are computed at the following (x/y) locations: (9/9), (15/9), (12/12), (9/15), and (15/15).

[0039] The classifier 445 can be trained using the classifier trainer 415 and tested using the classifier tester 420.

[0040] The training set 450 is a set of images that are used to train the classifier 445 using the classifier trainer 415. An image in the training set 450 is a synthetic image that was generated by using the renderer 410 to render a 3D model 460.

[0041] A 3D model 460 is a model of a three-dimensional object and can be rendered using the renderer 410 to produce a synthetic image.

[0042] The operating system tools 430 include a random number generator 455. The random number generator 455 can generate a random (or pseudo-random) number.

[0043] The renderer 410 is a conventional software application that can render a 3D model 460, such as Blender (an open source software package for 3D modeling and rendering). The renderer 410 renders a model 460 based on a set of rendering parameter values. In the example introduced above, there are six rendering parameters (the viewpoint's location in azimuth and elevation, the viewpoint's rotation around its optical axis, the point light source's location in azimuth and elevation, and the intensity ratio between the ambient light and the point light source).

[0044] In one embodiment, the renderer 410 renders a 3D model 460 at a resolution of 100x100 pixels. The renderer 410 also scales and smoothes the rendered (synthetic) image, resulting in a 23x23 pixel gray-value (synthetic) image. FIG. 5 shows examples of the rendered (synthetic) images after scaling and smoothing, according to one embodiment of the invention.

[0045] The classifier trainer 415 is a conventional software application that can train a classifier 445 given a training set 450 (e.g., a set of training images).

[0046] The classifier tester 420 is a conventional software application that can test a classifier 445 to determine the classifier's classification accuracy. In one embodiment, the classifier 445 is tested using 40,000 views per class, where each view is randomly drawn from the low-dimensional space of rendering parameters. The classifier's output is computed, using the real-valued output of the SVM. (For samples of the negative class (label "-1"), the output is multiplied by -1.) The most difficult views from each class are determined (e.g., the 100 views with the lowest accuracy rates).

[0047] The active learning module 405 includes a control module 435 and a local minima finder 440. The control module 435 controls the operation of the active learning module 405 so that the active learning module 405 can train a compact classifier 445 for view-based object recognition. The control module 435 is further described below with reference to FIG. 1.

[0048] The local minima finder 440 finds local minima (and the associated rendering parameter values) of a classifier's output in a low-dimensional rendering space. A set of views (e.g., the 100 most difficult views from each class) are used as starting points of an

optimization algorithm (e.g., the Nelder-Mead simplex algorithm). Note that the values of the six rendering parameters are already known for every view. A number of iterations of the algorithm (e.g., ten) are computed in order to find local minima (and the associated rendering parameter values) of the classifier's output in the rendering space. The classifier's output is computed, using the real-valued output of the SVM. (For samples of the negative class (label "-1"), the output is multiplied by -1.) FIG. 6 shows examples of the initially-selected (i.e., starting point) views and the views at the nearby local minima, according to one embodiment of the invention. In FIG. 6, example pairs are vertically arranged with the initial views on top and the views at nearby local minima on the bottom.

[0049] In general, active learning achieves the same error rate (e.g., equal error rate or "EER") with significantly smaller training sets and significantly fewer support vectors than training on a random selection of object views. In other words, active learning trains a "compact" classifier.

[0050] Returning to FIG. 1, in step 110, an initial training set 450 is generated, where the initial training set 450 includes one or more synthetic images. Each synthetic image is generated by the renderer 410 based on a 3D model 460 and a set of rendering parameter values. In one embodiment, the initial training set 450 is comprised of randomly selected samples from the low-dimensional space of rendering parameters.

[0051] In one embodiment, the control module 435 performs step 110. For example, the control module 435 uses the random number generator 455 to obtain one randomly-selected value for each of the rendering parameter values. (In the example introduced above, there are six rendering parameters, so each randomly-selected sample has one randomly-selected value for each of these six parameters.) The control module 435 then uses the renderer 410 to generate a synthetic image of a 3D model 460 based on the rendering parameter values. The synthetic image is then added to the initial training set 450. In one embodiment, the initial training set 450 includes 200 samples per class (e.g., 200 synthetic images per 3D model 460).

[0052] In step 120, the classifier 445 is trained using the training set 450. In one embodiment, the control module 435 performs step 120. For example, the control module 435 uses the classifier trainer 415 to train the classifier 445 using the training set 450. The first time that step 120 is executed, the training set 450 is the initial training set. The second and subsequent

times that step 120 is executed, the training set 450 will have been modified, as explained below in conjunction with step 150.

[0053] In step 130, local minima are found of the classifier's output in the low-dimensional rendering space. In one embodiment, the control module 435 performs step 130. For example, the control module 435 first uses the classifier tester 420 to test the classifier 445 and determine the classifier's classification accuracy. The classifier tester 420 determines the most difficult views from each class (e.g., the 100 views with the lowest accuracy rates). The control module 435 then uses the local minima finder 440 to find local minima (and the associated rendering parameter values) of the classifier's output, using the most difficult views as starting points.

[0054] In step 140, images are rendered at the local minima (using the associated rendering parameter values). In one embodiment, the control module 435 performs step 140. For example, the control module 435 uses the renderer 410 to render object views (synthetic images) at the local minima that were found in step 130. In one embodiment, 200 object views are rendered. The control module 435 also computes the orientation histograms of the rendered views. The orientation histograms will be used as the object features (i.e., the basis for classification).

[0055] In step 150, the newly-rendered images are added to the training set. In one embodiment, the control module 435 performs step 150. For example, the control module 435 adds the newly-rendered images (generated in step 140) to the training set 450 by storing the images appropriately. This generates a modified training set 450 that includes both the images that were rendered in step 140 and the images that were used to train the classifier in step 120.

[0056] In step 160, the procedure is repeated starting from step 120, so that the classifier 445 is retrained using the modified training set 450.

[0057] Although the invention has been described in considerable detail with reference to certain embodiments thereof, other embodiments are possible as will be understood to those skilled in the art. For example, another embodiment is described in "Object Recognition with 3D Models" by B. Heisele, G. Kim, and A. Meyer, Proceedings of the 2009 British Machine Vision Conference (BMVC), London, England, September 7-10, which is hereby incorporated by reference.

What is claimed is:

1. A computer-implemented method for training a view-based object recognition classifier, comprising:

generating an initial training set of images, wherein an image in the initial training set is generated based on a three-dimensional model and a set of one or more rendering parameter values;

training the classifier using the initial training set;

determining the classifier's accuracy;

determining a set of one or more local minima of the classifier's output;

for each local minimum in the set of local minima:

determining a set of one or more rendering parameter values associated with the local minimum; and

generating an additional image based on the three-dimensional model and the determined set of one or more rendering parameter values;

training the classifier using the initial training set and the generated additional images.

2. The method of claim 1, wherein the classifier comprises a statistical classifier.

3. The method of claim 1, wherein the classifier comprises a support vector machine.

4. The method of claim 1, wherein the classifier classifies an image based on a feature vector of the image's gradient histograms.

5. The method of claim 1, wherein a rendering parameter comprises one element of a group containing: an azimuth location of a viewpoint, an elevation location of a viewpoint, a rotation of a viewpoint around its optical axis, an azimuth location of a point light source, an elevation location of a point light source, and an intensity ratio between an ambient light and a point light source.

6. The method of claim 1, wherein a value for a rendering parameter for an image in the initial training set is determined randomly or pseudo-randomly.

7. The method of claim 1, wherein an image in the initial training set is associated with a label that specifies whether the image is a positive example or a negative example.

8. The method of claim 1, wherein determining the set of one or more local minima of the classifier's output comprises performing an optimization algorithm.

9. A machine-readable storage medium encoded with instructions that, when executed by a processor, cause the processor to perform a method for training a view-based object recognition classifier, comprising:

generating an initial training set of images, wherein an image in the initial training set is generated based on a three-dimensional model and a set of one or more rendering parameter values;

training the classifier using the initial training set;

determining the classifier's accuracy;

determining a set of one or more local minima of the classifier's output;

for each local minimum in the set of local minima:

determining a set of one or more rendering parameter values associated with the local minimum; and

generating an additional image based on the three-dimensional model and the determined set of one or more rendering parameter values;

training the classifier using the initial training set and the generated additional images.

10. The medium of claim 9, wherein the classifier comprises a statistical classifier.

11. The medium of claim 9, wherein the classifier comprises a support vector machine.

12. The medium of claim 9, wherein the classifier classifies an image based on a feature vector of the image's gradient histograms.

13. The medium of claim 9, wherein a rendering parameter comprises one element of a group containing: an azimuth location of a viewpoint, an elevation location of a viewpoint, a rotation of a viewpoint around its optical axis, an azimuth location of a point light source, an elevation location of a point light source, and an intensity ratio between an ambient light and a point light source.

14. The medium of claim 9, wherein a value for a rendering parameter for an image in the initial training set is determined randomly or pseudo-randomly.

15. The medium of claim 9, wherein an image in the initial training set is associated with a label that specifies whether the image is a positive example or a negative example.

16. The medium of claim 9, wherein determining the set of one or more local minima of the classifier's output comprises performing an optimization algorithm.

17. A system for training a view-based object recognition classifier, comprising:
- a machine-readable storage medium encoded with machine-readable instructions for performing a method, the method comprising:
 - generating an initial training set of images, wherein an image in the initial training set is generated based on a three-dimensional model and a set of one or more rendering parameter values;
 - training the classifier using the initial training set;
 - determining the classifier's accuracy;
 - determining a set of one or more local minima of the classifier's output;
 - for each local minimum in the set of local minima:
 - determining a set of one or more rendering parameter values associated with the local minimum; and
 - generating an additional image based on the three-dimensional model and the determined set of one or more rendering parameter values;
 - training the classifier using the initial training set and the generated additional images; and
 - a processor configured to execute the machine-readable instructions encoded on the machine-readable storage medium.
18. The system of claim 17, wherein a rendering parameter comprises one element of a group containing: an azimuth location of a viewpoint, an elevation location of a viewpoint, a rotation of a viewpoint around its optical axis, an azimuth location of a point light source, an elevation location of a point light source, and an intensity ratio between an ambient light and a point light source.
19. The system of claim 17, wherein a value for a rendering parameter for an image in the initial training set is determined randomly or pseudo-randomly.
20. The system of claim 17, wherein an image in the initial training set is associated with a label that specifies whether the image is a positive example or a negative example.

1/3




FIG. 1

2/3




FIG. 2




FIG. 3




FIG. 5




FIG. 6




FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 10/40678

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06K 9/62 (2010.01)

USPC - 382/159

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)

IPC(8): G06K 9/62 (2010.01)

USPC: 382/159

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC: 382/224; 345/419; 706/13, 25, 924 (keyword limited; terms below)

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

PubWest; Google Scholar; Google Patents; FreePatentsOnline; Dialog. Search terms used:

machine-vision MV machine-imaging automated-inspection optical-character-recognition OCR, training image-training training-set, classify classifier, recognize recognition, accurate accuracy image-recognition-quality, processor ...

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2006/0039593 A1 (SAMMAK et al.) 23 February 2006 (23.02.2006) entire document, especially Abstract; para [0007], [0014], [0053], [0060], [0082], [0099], [0100], [0116], [0119], [0125], [0167], [0185], [0237]	1 - 20
Y	US 2007/0031028 A1 (VETTER et al.) 08 February 2007 (08.02.2007) entire document, especially Abstract; Figs.11, 12; para [0003], [0024], [0029], [0048], [0068], [0083]	1 - 20
A	US 2006/0083423 A1 (BROWN et al.) 20 April 2006 (20.04.2006) entire document	1 - 20
A	US 2007/0276776 A1 (SAGHER et al.) 29 November 2007 (29.11.2007) entire document	1 - 20

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

* 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

13 August 2010 (13.08.2010)

Date of mailing of the international search report

23 AUG 2010

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

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