

IMAGE CLASSIFICATION AND LABELING

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims priority under 35 U.S.C. Section 119(e) to U.S. Provisional Patent Application No. 62/289,902, filed February 1, 2016, and titled "Method for Training an Algorithm to Classify Elements and Objects in Images," the entire content of which is incorporated herein by reference.

TECHNICAL FIELD

[0002] The present disclosure relates generally to image classification, and more particularly to training image classification models based on training images associated with multiple classes and classifying images into multiple classes using the trained image classification models.

BACKGROUND

[0003] Trademarks, logos, other graphically designed images, and other brand related intellectual properties are important business assets. Many countries provide formal mechanisms for people and companies to protect their IP, including trademarks, designs, etc. that may include graphically designed images, via trademark registration and design patent registration. For example, trademark registration systems may enable a company to register its name and logos in a trademark database and to object to or enforce against other names or logos that may be similar. Similarly, design patents in the U.S. and registration of industrial designs in other countries or regions (e.g., Europe) provide formal protection for designs that have a wide variety of graphical representation including line drawings, 3D CAD models along with 2D perspective views, 2D images like graphical user interfaces and surface ornamentation, etc.

[0004] Many trademark offices around the world label trademarks with a list of keywords or codes (e.g., design codes in the U.S., Vienna Codes in Europe, and image constituents in Australia) that semantically describe elements or objects inside, for example, a graphical image. To illustrate, U.S. design codes and Vienna Codes are class labels that are semantically descriptive of images and/or objects in images. These class labels have hierarchical structures and are used to classify the figurative elements of an image. In general, the hierarchical structures of labels used in some IP registration systems proceed from the general to the more

- 2 -

particular, for example, considering the figurative elements of trademarks in hierarchical levels starting from categories, followed by divisions that are followed by sections. Some IP registration systems may use single level (non-hierarchical) systems while others may have less or more than three levels of hierarchy.

[0005] For design patents and industrial designs, many IP offices have adopted a classification system called Locarno classifications and product indications to both describe a product and its appearance. The U.S. has its own set of design classifications called U.S. Design Classes. Both Locarno and U.S. Design Classes are hierarchical in nature with the top level typically describing the function of the product and sub classes describing specific functional feature, distinctive ornamental appearance, or form. Non-governmental organizations may also use similar types of labels for classifying and/or annotating the graphical assets such as logos and designs.

[0006] IP owners have to enforce their IP rights through watching and addressing potential infringement and efforts to register similar trademarks and designs. However, finding infringing use of images, for instance, on e-commerce sites, or finding misuses of a brand in social networks or the internet can be challenging. Further, companies have challenges in searching for registered trademarks or designs to determine whether to use and/or attempt to register a name, logo, or design.

[0007] In many IP registration systems including in the U.S. and Europe, multiple labels may be associated with an image or with an element/object in an image (e.g., a trademark or a design). The labels may also be hierarchical labels. Labels that may be associated with an image or an image object/element may be herein referred to as multi-labels.

[0008] Thus, systems and methods that enable classifying images and/or objects/elements in images into multiple classes are desirable.

SUMMARY

[0009] The present disclosure relates generally to image classification, and more particularly to training image classification models based on training images associated with multiple classes (with class labels) and classifying images into multiple classes (with class labels) using the trained image classification models. In an example embodiment, a method of training an image classification model includes obtaining training images associated with labels,

where one or more training images of the training images are each associated with two or more labels of the labels and where one or more labels of the labels each correspond to an image classification class. The method further includes classifying training images into one or more classes using a deep convolutional neural network, and comparing the classification of the training images against labels associated with the training images. The method also includes updating parameters of the deep convolutional neural network based on the comparison of the classification of the training images against the labels associated with the training images.

[0010] In another example embodiment, a method of classifying images using one or more image classification models includes obtaining training images associated with labels, wherein each training image of the training images is associated with two or more labels of the labels and wherein each label of the labels corresponds to an image classification class and wherein the labels have a hierarchical structure having one or more hierarchy levels. The method further includes training a separate deep convolutional neural network for each hierarchy level of the labels using the training images and the labels associated with the training images. The method also includes classifying an input image into one or more classes based on the trained deep convolutional neural network for each hierarchy level, wherein classifying the input image into the one or more classes depends on prior probabilities of accurately classifying other images into the one or more classes.

[0011] In yet another example embodiment, a system for classifying images using one or more image classification models includes a training image source comprising training images associated with labels, wherein the labels include codes used by a trademark registration organization or an industrial design registration organization. The system further includes an image processing device communicably coupled to the training image source, wherein the image processing device is configured to obtain the training images from the training image source and to train a deep convolutional neural network using the training images and the labels associated with the training images. The system also includes an input image source comprising input images to be classified into image classification classes, wherein classifying an input image into one or more image classification classes associates the input image with one or more labels corresponding to the one or more image classification classes.

[0012] These and other aspects, objects, features, and embodiments will be apparent from the following description and the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] Reference will now be made to the accompanying drawings, which are not necessarily drawn to scale, and wherein:

- 4 -

[0014] FIG. 1 illustrates a system for training image classification models and for classifying images according to an example embodiment;

[0015] FIG. 2 illustrates a method of training image classification models and classifying images according to an example embodiment;

[0016] FIG. 3 illustrates a method of training image classification models (i.e., convolutional neural networks) according to an example embodiment;

[0017] FIG. 4 illustrates a method of classifying images using trained classification models according to an example embodiment;

[0018] FIG. 5 illustrates classification/labelling of an image using the system of FIG. 1 and based on the methods of FIGS. 2, 3, and 4; and

[0019] FIGS. 6-8 illustrate classification of input images based on trained classification models for each hierarchy level of the multi-labels according to an example embodiment.

[0020] The drawings illustrate only example embodiments and are therefore not to be considered limiting in scope. The elements and features shown in the drawings are not necessarily to scale, emphasis instead being placed upon clearly illustrating the principles of the example embodiments. Additionally, certain dimensions or placements may be exaggerated to help visually convey such principles. In the drawings, reference numerals designate like or corresponding, but not necessarily identical, elements.

DETAILED DESCRIPTION OF THE EXAMPLE EMBODIMENTS

[0021] In the following paragraphs, example embodiments will be described in further detail with reference to the figures. In the description, well known components, methods, and/or processing techniques are omitted or briefly described. Furthermore, reference to various feature(s) of the embodiments is not to suggest that all embodiments must include the referenced feature(s).

[0022] In some example embodiments, systems and methods that are based on convolutional neural network may be used to classify images into multiple classes such that the images are each associated with respective multiple labels. To illustrate, one or more convolution neural networks (also referred to herein as a classification models) are first trained, and the trained convolutional neural networks are then used to classify images. Labels that can

be associated with a single image or with an element/object in an image may be referred to herein as labels, multiple labels, or multi-labels.

[0023] In some example embodiments, during a training operation, a convolutional neural network (e.g., a deep convolutional neural network) may be trained such that parameters of the network are learned using training images tagged (for example, manually) with multi-labels that describe visual-semantic meanings of the images. During classification operations, the trained convolutional neural network is used to classify other images such that multi-labels are associated with the images. For example, the images classified by the trained network may be tagged or otherwise associated with multi-labels that describe visual-semantic meanings of the images. As described in more detail below, the trained convolutional neural network (i.e., the trained classification model) may include several layers including convolutional, pooling, activation, gating, dense connection and drop out to robustly classify an image such that the image is associated with appropriate multi-labels.

[0024] In some example embodiments, the multi-labels associated with training images and also used to classify images may be hierarchical labels. To illustrate, when the multi-labels are hierarchical labels, a respective trained classification model (i.e., a respective convolutional neural network) may be trained for each level of the hierarchical labels. For example, three convolutional neural networks may be trained when the training images are tagged or otherwise associated with multi-labels that have three hierarchies. During classification operations, multiple trained classification models may be used to classify images such that the images are tagged or otherwise associated with the hierarchical multi-labels at each hierarchy level.

[0025] To illustrate, images and associated semantic descriptive labels, such as U.S. design codes, Vienna codes, and keywords image constituents, used in trademark databases can be used to train very large deep convolution neural networks that can encode the semantic information into feature maps. The trained convolution neural networks (i.e., the classification models resulting from the training operation) may then be used to classify other images such that the images are associated with respective semantic descriptive labels used in the trademark databases.

[0026] For example, Table 1 illustrates example labels that are based on semantic descriptions of images that may be registered as trademarks in the U.S. Please note that the

particular images that have been tagged or otherwise associated with the labels are not shown in Table 1.

Image	Labels using U.S. Design Codes		
	"Category"	"Division"	"Section"
Image 1	24 - Heraldry, flags, crowns, crosses, arrows and symbols	24.05- Seals	24.05.01 - Seals, circular or elliptical; Circular or elliptical seals
	26 - Geometric figures and solids	26.01- Circles 26.05- Triangles	26.01.11 - Circles comprised of plants; Letters, numerals, punctuation, geometric figures, objects, humans, plants or animals comprising a circle; Circles comprised of punctuation 26.05.13 - Triangles, exactly two triangles; Two triangles 26.05.16 - Triangles touching or intersecting 26.05.21 - Triangles that are completely or partially shaded
Image 2	01 - Celestial bodies, natural phenomena, geographical maps 18 - Transport; equipment for animals; traffic signs 26 - Geometric figures and solids	01.07 -- Globes 18.09 - Air or space vehicles 26.17 - Lines, bands, bars, chevrons and angles	01.07.25 - Globes, other 18.09.01 - Gliders, airplane; Airplanes 26.17.07 - Speed, lines depicting; Propulsion, lines depicting; Lines depicting speed, propulsion, heat or wind; Wind, lines depicting; Heat, lines depicting

[0027] As illustrated in Table 1, the multi-labels are intended to provide semantic descriptive information about the image 1 and image 2. The labels in Table 1 also have a hierarchical structure where the hierarchies are "Category," "Division," and "Section." For example, the highest hierarchy, "Category," provides more general semantic description of the images compared to the lower hierarchies. The next highest hierarchy, "Division," provides

more general semantic description of the images compared to the lower hierarchy, "Division," but more specific semantic description than the highest hierarchy, "Category." In some alternative embodiments, hierarchical multi-labels may have a different hierarchical relationship than shown in Table 1 without departing from the scope of this disclosure. In contrast to some image databases/datasets, each image in Table 1 is associated with multiple labels at each hierarchy.

[0028] FIG. 1 illustrates a system 100 for training image classification models and for classifying images according to an example embodiment. The system 100 may include an image processing device 102, a training image source 104, an input image source 106. The image processing device 102 may communicate with the training image source 104 and the input image source 106 over a network 108. For example, the network 108 may be the internet, a local area network, a wide area network, or a combination thereof. Alternatively or in addition, the network 108 may be or may include a cable connection such as, for example, a universal serial bus (USB) cable connection.

[0029] In some example embodiments, the training image source 104 may include one or more databases of images tagged or otherwise associated with multi-labels. The images from the training image source 104 may be used by the image processing device 102 to train convolutional neural networks (i.e., classification models) that can be used to classify other images such as the images from the input image source 106. The training images may be graphically designed images (e.g., line drawings, 2D perspective views of 3D CAD models, 2D drawings, etc.), photos, other types of images, or a combination of two or more of the preceding. The training images may be trademarks, logos, and product designs including images used in design patents and industrial design registrations, etc. The training image source 104 may include image databases that include millions of images or few images depending on a particular application. The training image source 104 may be one or more standalone data storage devices, data storage devices of a computer or computer network, or another device or system that may be used to provide images along with the associated multi-labels. For example, the training image source 104 may include a storage device (e.g., a static random access memory or another type of computer readable medium). The multi-labels associated with individual training images provided by the training image source 104 to the image processing device 102 may be in a form of metadata of the individual training images or in another format that otherwise associates the

individual training images with respective multi-labels. In some example embodiments, hundreds or thousands of possible visually semantic (i.e., semantically descriptive) labels may be potentially used to label training images. The training image source 104 may be a storage device (e.g., a static random access memory, etc.) that is communicably coupled to the image processing device 102.

[0030] In some example embodiments, the multi-labels associated with the images from the training image source 104 may have a hierarchical structure as described above. Alternatively, the multi-labels may be non-hierarchical. As a non-limiting example, the training image source 104 may include a database of one or more IP registration organizations, such as the U.S. Patent and Trademark Office, and the multi-labels may be semantically descriptive labels, such the U.S. Design Codes or other semantically descriptive codes.

[0031] In some example embodiments, the input image source 106 may include one or more databases of images that can be classified by the image processing device 102 using convolutional neural networks (i.e., classification models) that are trained using the training images that are tagged or otherwise associated multi-labels and provided by the training image source 104. The images included in the input image source 106 may be graphically designed images (e.g., line drawings, 2D perspective views of 3D CAD models, 2D drawings, etc.), photos, other types of images, or a combination of two or more of the preceding. The images provided by the input image source 106 may be trademarks, logos, product designs including images used in design patents and industrial design registrations, etc. The input image source 106 may include image databases that include millions of images or few images depending on a particular application, and may be classified into hundreds or thousands of classes (i.e., the images may be labeled by some of possible hundreds or thousands of labels). The input image source 106 may be one or more standalone data storage devices, data storage devices of a computer or computer network, or another device or system that may be used to provide images to be classified by the image processing device 102. For example, the input image source 106 may include a storage device (e.g., a static random access memory or another type of computer readable medium). In some example embodiments, the input image source 106 may include a user device that submits one or more images that a user want classified by the image processing device 102. Alternatively or in addition, the input image source 106 may include a web page that includes images with metadata.

[0032] In some example embodiments, the image processing device 102 includes a processing unit 110, such as a central processing unit, a memory device 112, a storage device 114, and a network interface 116 as well as other components that support the function of the image processing device 102 to train classification models and classify images. The network interface 116 may be used by the image processing device 102 to communicate with other devices including the training image source 104 and the input image source 106. For example, one or both of the training image source 104 and the input image source 106 may be local storage devices that are communicably coupled to the image processing device 102 by respective cable (e.g., a USB cable or CAT 5 cable) via the network interface 106. In some example embodiments, the training image source 104 and the input image source 106 may be a single device that includes training images with associated multi-labels as well as images that need to be classified.

[0033] In some example embodiments, the processing unit 110 may include multiple processor units including specialized (e.g., graphics) processing units. For example, computer executable code may be stored in the storage medium 112 and may be executed by the processing unit 110 to implement the training of classification models based on the training images and the classification of images based on the trained classification models. For example, the memory device 112 may include one or more static random access memory devices or another type of non-transitory computer readable media.

[0034] In some example embodiments, the storage device 114 may be used to store executable code and/or data including image data, labels, and/or other information related to the training of classification models and the classification of images. For example, in some embodiments, the storage device 114 may be used to store training images previously received from the training image source 104 or another source of training images that have associated labels. Alternatively or in addition, the storage device 114 may be used to store images to be classified by the image processing device 102. The storage device 114 may include one or more static random access memory devices or another type of non-transitory computer readable media.

[0035] In some example embodiments, the memory device 112 and/or the storage device 114 may be used to store a list of multi-labels applicable to a particular database of training images. For example, a complete or a particular portion of semantically descriptive labels used by an organization to classify trademark or design images (e.g., graphically designed images)

may be stored in the storage device 114 for use by the image processing device 102 during the training of the classification models and/or during the classification of images.

[0036] By using convolutional neural network(s) trained with training images and multiple labels associated with individual images of the training images, the system 100 can perform reliable classification of images into multiple classes. The use of the convolutional neural networks described herein to classify images improves the image classification capability of the image processing device 102 by enabling classification of images into multiple classes (i.e., associate images with multiple labels) as compared to systems and neural networks that are limited to mutually exclusive labels, where each image (or set of pixels in an image) is associated with a single label and classified accordingly. For example, in systems and neural networks that are limited to mutually exclusive labels, an image can be considered to contain “car” or “human” labels but not both. In contrast, the labels used by the system 100 may be non-mutually exclusive labels, where a label is may be associated with multiple images or objects in images. Further, the system 100 enables labels to be associated with multiple classes of images. For example, if a class of images is that of dogs and another class of images is that of cats, the system 100 enables a semantically descriptive label (e.g., furry) to be associated with both images from both classes.

[0037] In some example embodiments, the image processing device 102 may be implemented using hardware (e.g., FPGA or a microprocessor), software, or a combination of hardware and software. Further, some of the components of the image processing device 102 may be omitted or integrated into a single component without departing from the scope of this disclosure. Although the image processing device 102 is shown as a single device, in some example embodiments, the image processing device 102 may be a system of multiple components/devices. Although particular components of the system 100 are shown in FIG. 1, the system 100 may include fewer or more components than shown. Descriptions provided herein with respect to an image or images may be applicable to objects/elements in the image(s).

[0038] FIG. 2 illustrates a method 200 of training image classification models and classifying images according to an example embodiment. Referring to FIGS. 1 and 2, in some example embodiments, the method 200 includes obtaining training images with associated labels at step 202. For example, the image processing device 102 may obtain training images tagged or otherwise associated with multi-labels from the training image source 104. Alternatively, the

image processing device 102 may obtain the training images and associated labels from another device such as the storage device 114.

[0039] The method 200 includes training one or more classification models (i.e., convolutional neural networks) at step 204. For example, the system 100 of FIG. 1 may be used to train one or more convolutional neural networks. The fundamental concepts of convolution layers are weight sharing and encoding multiple feature maps. Weight sharing allows extracting visual information regardless spatial location (translational invariance) and multiple parallel feature maps allow extracting different types/levels of details of an image simultaneously. Description of convolutional neural networks may be found in Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," Proc. IEEE, vol. 86, no. 11, pp. 2278–2323, 1998, the content of which is incorporated herein by reference.

[0040] As described in more detail with respect to FIG. 3, a convolutional neural network trained based on the training images obtained at step 200 may be a deep convolutional neural network that includes several layers. The training of a convolutional neural network includes tuning/updating of weights and biases of the network kernels by iteratively doing gradient optimization on labelled data (e.g., human labelled data). In contrast to convolutional neural networks that use softmax activation or sigmoid activation at the classification layer, the convolutional neural networks used by the system 100 and in the method 200 use soft-sigmoid activation described below for reliable classification of images into multiple classes (i.e., to associate multiple images with an image). For example, convolutional neural networks that use softmax activation at the classification layer require mutually exclusive classes/labels. Convolutional neural networks that use sigmoid activation at the classification layer or sigmoid activation of the last layer or sigmoid-cross entropy for the objective function during training can suffer from negative saturation and hanging within local minima. By using the convolutional neural network with soft-sigmoid activation at the classification layer or soft-sigmoid-cross entropy for the objective function during training, the system 100 and the image processing device 102 may train one or more convolutional neural networks using images tagged or otherwise associated with multiple labels (i.e., multi-labels) and may reliably and efficiently classify images into multiple classes (i.e., associate images with multiple labels) using the trained one or more convolutional neural networks. In some alternative embodiments, the system 100

and the method 200 may use a different means to prevent or recover from negative saturation without departing from the scope of this disclosure.

[0041] In some example embodiments, a respective convolutional neural network may be trained based on each hierarchical level of hierarchical multi-labels (e.g., U.S. Design codes, Vienna codes, etc.) associated with images or objects/elements in images. In general, the training images and associated multi-labels used in the training of convolutional neural network(s) used by the system 100 and in the method 200 may be non-hierarchical or hierarchical.

[0042] In some example embodiments, the training images may be optionally pre-processed, for example, for better data augmentation. For example, the training images may be resized to fixed height and width. The aspect ratio of the images may be maintained by padding the shorter size with, for example, black, white and edge color. Grayscale and binary version of RGB images may also be used as data augmentation. As another example of pre-processing, the training images may be pre-whitened, and random contrast, random brightness and random flips and rotation may also be applied during each iteration.

[0043] In some example embodiments, the method 200 may include testing the trained convolutional neural network(s) (i.e., trained classification models) at step 206. For example, a sample of the training images and associated multi-labels from the training image source 104 or another source may be provided to the image processing device 102 to test/benchmark the trained convolutional neural network(s). Additional refinement of the convolutional neural network(s) may be performed, for example, using more training images if the performance of the system 100 or the method 200 is below a threshold. In some alternative embodiments, the testing of the trained classification models may be omitted without departing from the scope of this disclosure.

[0044] In some example embodiments, the method 200 includes classifying input images using the one or more trained classification models (i.e., trained convolutional neural network(s)) at step 208. For example, the input images may be received by the image processing device 102 from the input image source 106 or from another source such as the storage device 114. The image processing device 102 may classify the input images by based on one or more trained classification models. For example, for a particular dataset of images that are to be classified based on non-hierarchical labels, the system 100 may train a single classification model that is

used to classify corresponding input images. For another dataset of images that are to be classified based on hierarchical labels, the system 100 may train a single classification model that is used to classify corresponding input images. By classifying the input images into multiple classes, the image processing device 102 effectively labels, tags or otherwise associates classified input images with multi-labels. In some example embodiments, one or more of the input images may be classified in single class when the images or objects/elements in the images fall only under a single class.

[0045] FIG. 3 illustrates a method/system 300 of training image classification models (i.e., convolutional neural networks) according to an example embodiment. Referring to FIGS. 1-3, a deep convolutional neural network may include convolution and pooling layers 306, fully connected layers 310, and a logits layers 312 as illustrated in FIG. 3. In the convolution layers, weight sharing kernels are scanned on input feature-maps (image or output of the previous layers), and non-linear activation function like ReLU, ELU, sigmoid and etc. are used to generate input for the next set of layers. Usage of weight sharing kernels enables translational invariance and multiple parallel kernels enable extracting different visual information.

[0046] A pooling layer is used in between convolution layers to reduce the dimension of input feature map. In pooling, weight-sharing kernel is scanned with more than one-pixel stride length (commonly stride length of two is used) and get the maximum value within the kernel region. Instead of maximum, any statistical aggregation function can be used (maximum: max-pooling, min: min-pooling and etc.). Pooling layer helps to reduce the dimensionality and disregard noise.

[0047] These convolution layers and pooling layers of the convolution and pooling layers 306 are repeated number of times to make sure visual information are extracted in multiple levels (from global details to finer details). When the network is deeper with multiple convolution and pooling layers, image information from beginning of the network is transferred to the deeper side of the network using gating 308.

[0048] Hyper parameters of the network (number of layers, kernel size, number of gating and number of kernels) are set based on complexity of image data and number of labels used in application. Part of the network also can be repeated in parallel and concatenated together if an application is more complex and requires learning more feature maps to different classification labels, resulting in the network being relatively wide.

[0049] In general, the convolution and pooling layers 306 perform the following main operations: 1) Convolve weight sharing kernels on the pre-processed input image; 2) Apply relevant pooling (max, mean, min) on the generated feature maps to accommodate image invariances (e.g., translation); and 3) Use an appropriate activation function (ReLU, Sigmoid, ELU, etc.) to map the pooled feature to non-linear output. These main operations are repeated with a combination of different kernel sizes (3x3, 5x5, 11x11), activation functions and gating (highway, skip) to capture the different levels of visual features. With deeper convolution network, smaller kernel size with larger number of feature maps are used to adapt smaller receptive field with finer details.

[0050] The outputs of the last convolution layers of the convolution and pooling layers 306 are concatenated together and densely mapped to fully connected layers 310. Different scales of drop-out are used to make sure the network is not over-fitted. The fully connected layers 310 make sure the visual information extracted from convolution layers are narrowed to map classification labels.

[0051] The logits layer (i.e., soft-sigmoid layer) 312 is used to map the output of last fully connected layer (H) to get the confidence scores for each classification labels. The soft-sigmoid layer is defined as:

$$\text{soft-sigmoid} : H \mapsto \gamma * F_g(H) + (1 - \gamma) * F_s(H)$$

[0052] Where F_s is a sigmoid function. F_g is any non-linear function that satisfies $\partial F_g(H)/\partial H \neq 0$, (function that has non-zero gradient value for all the range of input). An example function that works in that context is softmax. Other activation functions like ReLU, ELU also satisfy this requirement. γ sets the weights to F_g and F_s . During the training of the convolutional neural network, γ is set to a function that correlates with cumulated cross-entropy loss (boosting F_g when the model is with higher loss). This process helps to bring back the false positive and false negative confidence scores that hangs in the negative saturated regions of sigmoid. During classification of input images using the trained convolutional neural network, γ is set to 0 to produce independent confident scores for each label.

[0053] Training the convolutional neural network using the images 302 with multi-labels 304 is done by back-propagating the network to minimize the cross-entropy between soft-sigmoid activation logits and the ground truth labels. Updating of parameters such as weights

and biases and other optimizations are performed by the optimization and updating module 316 based on output of the soft-sigmoid cross-entropy module 314. Mini-batch with the size of (64, 128, 256) from random shuffled input image data queue may be used as input to generalize loss function between the training images that belong to different classes.

[0054] As described above, when the labels 304 have a hierarchical structure, separate models are trained for each level of the hierarchy of labels. The number of convolution layers, kernel size, type of activation, pooling and gating, number of fully connected layers, number of neurons on the fully connected layers varies with the number of labels, number of training images used, type of labels used and type of training images etc. For example, if the network needs to classify large number of labels (thousands or more) a convolutional neural network with a large number of layers (similar to vgg16, Resnet101) is more appropriate to capture the complex structure of the visual features.

[0055] FIG. 4 illustrates a method/system 400 of classifying images using trained classification models according to an example embodiment. Referring to FIGS. 1-4, the trained classification model (i.e., trained convolutional neural network) that resulted from training of the classification model described above includes convolution and pooling layers 406, fully connected layers 410, and a logits layer 412 as illustrated in FIG. 4. The trained convolutional neural network of FIG. 4 is effectively the convolutional neural network of FIG. 3 at the end of the training process. To illustrate, the convolution and pooling layers 406 corresponds to the convolutional neural network 306, the fully connected layers 410 corresponds to the fully connected layers 310 and the logits layers 412 corresponds to the logits layers 312 with the exception of difference in parameter settings. Input images 402 that need to be classified may be provided from an image source such as the input image source 106 or another source. In some example embodiments, the input images may be pre-processed in a similar manner as described with respect to the training images.

[0056] For a given input image, the trained classification models from the step 204 of the method 200 and from the method of 300 output confident scores for each label. The output of the soft-sigmoid activation logits layer 412 generate independent sigmoid confidence score for each label in contrast to convolutional neural networks that use softmax activation resulting in confident scores that are distributed between all labels to have a summation of the score equals to one. Confidence scores indicate the level of confidence that an object in the image or an

image is classified correctly in a particular class (i.e., a particular label). For example, a score of 1 may indicate a highest confidence level, and a score of 0 may indicate the lowest confidence level. By using soft-sigmoid activation, the methods 200, 300, and 400 enable classification of an image and objects/elements in an image into multiple classes as shown by the classification/labeling output 414. That is, multiple labels (i.e., multi-labels) may be associated with an image. For example, an image may be tagged with multiple labels and may not be limited to a single class.

[0057] FIG. 5 illustrates classification/labelling of an image (e.g., a graphically designed image) using the system 100 and based on the methods 200, 300, and 400. As illustrated in FIG. 5, the image 502 is classified into three classes (i.e., associated with three multi-labels) as shown by the classes/labels 504 along with confidence scores 506 provided in percentages.

[0058] FIGS. 6-8 illustrate classification of input images based on trained classification models for each hierarchy level of the multi-labels according to an example embodiment. In some example embodiments, separate classification models may be trained for each level of hierarchy of hierarchical labels, as described above. To illustrate, trademark labels such as US-design codes and Vienna-codes are built upon a hierarchical structure from "Category" → "Division" → "Section," as illustratively shown in Table 1. The trained classification models for each hierarchy level (for example, H(C) for "Category", H(D) for "Division" and H(S) for Section) may be used to determine a final classification of an image or an object in an image. The hierarchy information together with information on each of the classification models from prior statistics may be used to robustly classify each category, division and section labels. FIGS. 6-8 illustrate how the identified labels for each of the hierarchical levels are used to refine the final scores.

[0059] To illustrate, probability scores for each class/label are computed from higher level to lower label in hierarchy. Confidence score for the higher level label from each model is computed and averaged as shown in FIG. 6. For each lower level, the probability score for the current level label is multiplied by the probability score of the corresponding higher level label as shown in FIGS. 7 and 8.

[0060] With the large amount of training images, prior statistics of each classification model's reliability towards correctly indicating a class/label to be associated with an input image may also be incorporated in computing the scores. To illustrate, prior probabilities that measure

the accuracy of each hierarchy level classification model (for example, $H(C)$ for “Category”, $H(D)$ for “Division” and $H(S)$ for Section) for a given class are pre-computed based on the statistics that extracted from, for example, millions of training data. For example, $P(C|c1)$ is how good the category model in predicting on $c1$ class. Confidence scores for a label, (e.g., $c1$) from its children’s model (D or S) is computed by multiplying the maximum confidence of its child labels ($D(c(di) == c1)$ or $S(c(si) == c1)$) from its relevant model with the prior statistics of the classification model’s confidence score for the given label. The final confidence score for the label can be defined by mean of confidences from each of its child models, multiplied by the parent probability. The parent probability (confidence score) is defined by the probability of the upper level hierarchy label to which that the child label belongs.

[0061] The above described systems and methods may be used to train deep convolutional neural network(s) to classify images based on semantically descriptive (semantic visual) labels that may have hierarchical structures. Example descriptions of three-level hierarchy labels used in US design codes to describe trademark images are shown in Table 1 above. Examples of multi-labels in the design patent/industrial design space include Locarno classifications and Product Indications.

[0062] As described above, in an example embodiment of a method of training an image classification system, a large-scale logo dataset may be used for training convolutional neural networks (i.e., classification models). The trademark images (typically graphically designed images) and the trademark image codes are used as labels as input into the deep learning system (i.e., a system based on deep convolutional neural network(s)). Use of trademarks as training image dataset provides a large resource of images that may be used to train deep convolutional neural networks for classification of images, logos, etc. As described above, the systems and methods described herein may be used in applications involving classification of industrial designs and a combination of trademarks and designs.

[0063] System benchmarking may be performed against both logo datasets (trademarks) and large scale photo datasets (ImageNet), as well as more application specific datasets such as e-commerce site images (usually product photos which can be either photos, digitally altered merchandise, or just plain logos). As described above, images, both training images and input images to be classified, may undergo preprocessing to normalize white balance, remove noise, scale images, etc.

[0064] Additionally, the metadata including labels and other relevant information can be preprocessed. Not all trademark meta data will be useful in general, some descriptors are more meaningful than others (for instance describing an image as containing segments is less useful than describing an image as 2 connected segments forming a checkmark). Relevance statistics for these labels can be collected beforehand to feed into the supervised training process of a deep network.

[0065] Additionally, as described above, different levels of metadata (i.e., hierarchical labels) can be differentiated, from most specific to least specific. For instance, a very specific meta data label for an image might be Trademark application number 11111 (trademark identified), a higher level label is the brand name "Twitter" (brand/owner), a higher level yet again is "Sparrow", "Bird", "Animal" (for instance, like the VC category, division, section hierarchy e.g. "02.01.01", "02.01", "02"). These different levels of description and labels are used for training.

[0066] The above formatted information may be used to train a convolutional neural network (i.e., classification model), which comprises of a multi-level architecture of alternating convolution and max pooling layers, typically with a final layer as a classifier, as described above. In an example embodiment, the above labels can be used to train the last layer. In systems and methods that are based on convolutional neural networks, convolution and max pooling reduces dimensionality through effectively clustering the features resulting from the convolution. Each layer of convolution provides a set of higher-level features with different representation. Understanding what architecture suits the level of semantic labels will help with training.

[0067] The systems and methods described herein enable accurate classification of visual elements or objects in graphically designed images such as trademarks, logos, and images of industrial designs including product designs. Taking advantage of hierarchical levels of multi-labels by training per-hierarchy classification models may result in even more reliable classifications. Objects and elements in photographs may also be classified using the systems and methods described herein. Applications include improving clearance search and watches for image-based trademarks (and common law logos) and industrial design for protection against infringement, as well as wide applicability in a variety of domains such as media monitoring, security, and consumer applications.

[0068] The systems and methods described herein may be used in many applications. For example, by using already established code standards for imaging labelling, the systems and methods may be used to protect intellectual property of graphically designed images, where the images are labelled by examiners according to the industry standard (Vienna code, design code and keyword, etc.). The labels can be used to train classification models on hierarchically labelled classes as mentioned above, and the trained classification models may be used to generate code suggestions for examiners during new registrations and help to find similar design/trademark images by providing the image or an object in the image desired to be registered to the classification system described herein. Such code generation capability may also be used with design patents and industrial design registrations. The trained classification models can also be used to look for infringement by similar design of a graphic image.

[0069] In some example embodiments, image searches may be performed based on the class-labels. To illustrate, combination of object labels describe the uniqueness of an. The trained classification models may be used to label or otherwise associate with labels images and objects in images. By using the trained classification models that can automatically generate all the combination of object labels, the image can be searched against large scale dataset for similarity.

[0070] In some example embodiments, web-image search may be performed based on metadata. Similar approaches as described above can be used to search for similar image on the web. Classification models trained on metadata also can be used to automatically index an uploaded image to the database. The systems and methods described herein may also be used to perform trademark and design clearance search and provide watch services.

[0071] In some example embodiments, reverse-image search may be performed using the systems and methods described herein with additional features from the outputs of the trained networks (visually encoded feature maps – intermediate network layers and/or auto generated semantic labels). Further, by having separate models for different hierarchy levels of labels, image search can be performed to look into different level of details in similarity. For example, matching twitter logo can be performed with the expectation of getting the results with: i) the exact shape/color/texture of the twitter logo; ii) a logo with the combination of birds in flight or with outspread wings, woodpeckers, other birds; pigeons; ravens; robins; crows; cardinals; doves; iii) type of animals; or iv) a logo with any animal.

[0072] In general, the systems and methods described herein can be used to train very large (wide/deep) deep learning models on the world's largest corpus of logo training image data with independent, hierarchical, multi label information. Such systems and methods can easily handle over 5 million trademark images and thousands of class labels during the training processes..

[0073] Although embodiments of the methods have been described, it is to be understood that such a system is not necessarily limited to the specific features or algorithms described. Rather, the specific features and algorithms are disclosed as example forms of implementing such systems.

[0074] Although some embodiments have been described herein in detail, the descriptions are by way of example. The features of the embodiments described herein are representative and, in alternative embodiments, certain features, elements, and/or steps may be added or omitted. Additionally, modifications to aspects of the embodiments described herein may be made by those skilled in the art without departing from the spirit and scope of the following claims, the scope of which are to be accorded the broadest interpretation so as to encompass modifications and equivalent structures.

CLAIMS

1. A computer implemented method of classifying images using one or more image classification models, the method comprising:

obtaining training images associated with labels, wherein each training image of the training images is associated with two or more labels of the labels and wherein each label of the labels corresponds to an image classification class, and wherein the labels have a hierarchical structure having one or more hierarchy levels;

training a separate deep convolutional neural network for each hierarchy level of the labels using the training images and the labels associated with the training images; and

classifying an input image into one or more classes based on the trained deep convolutional neural network for each hierarchy level, wherein classifying the input image into the one or more classes depends on prior probabilities of accurately classifying other images into the one or more classes.

2. The method of Claim 1, wherein a classification layer of the deep convolutional neural network is based on soft-sigmoid activation, wherein a logits layer of the deep convolutional neural network is used to map an output of a last fully connected layer of the deep convolutional neural network, wherein the logits layer is defined as $H \mapsto \gamma * F_g(H) + (1 - \gamma) * F_s(H)$, wherein H corresponds to the last fully connected layer, wherein F_g is any non-linear function that satisfies $\partial F_g(H)/\partial H \neq 0$, wherein F_s is a sigmoid function, and wherein γ sets weights for F_g and F_s .

3. The method of Claim 1, wherein the training images and the input images include graphically designed images.

4. The method of Claim 1, wherein the labels are non-mutually exclusive labels.

5. The method of Claim 1, wherein the labels are codes used by a trademark registration organization.

6. The method of Claim 1, wherein the labels are codes used to classify design patent images or industrial design images.

7. The method of Claim 1, wherein the labels are available as metadata of the training images associated with the labels.

8. The method of Claim 1, wherein classifying the input image into two or more classes based on the trained deep convolutional neural networks includes tagging or labelling the input image with two or more labels corresponding to the two or more classes.

9. The method of Claim 1, further comprising pre-processing the training images, wherein training the deep convolutional neural network using the training images and the labels associated with the training images comprises training the deep convolutional neural network using the pre-processed training images and the labels associated with the training images.

10. A computer implemented method of training an image classification model, the method comprising:

obtaining training images associated with labels, wherein each training image of the training images are associated with two or more labels of the labels and wherein each label of the labels corresponds to an image classification class, and wherein the labels have a hierarchical structure having one or more hierarchy levels;

classifying the training images into one or more classes using a separate deep convolutional neural network for each hierarchy level of the labels;

comparing the classification of the training images against labels associated with the training images; and

updating parameters of the deep convolutional neural network for each hierarchy level of the labels based on the comparison of the classification of the training images against the labels associated with the training images, wherein probabilities of accurately classifying other images into the one or more classes is pre-computed for use in classifying an input image into the one or more classes.

11. The method of Claim 10, wherein the training images include graphically designed images.

12. The method of Claim 10, wherein a separate deep convolutional neural network is trained for each hierarchy level of the labels.

13. The method of Claim 10, wherein the labels are codes used by a trademark registration organization.

14. The method of Claim 10, further comprising pre-processing the training images before classifying the training images comprises classifying the pre-processed training images.

15. The method of Claim 10, wherein a classification layer of the deep convolutional neural network is based on soft-sigmoid activation.

16. A system for classifying images using one or more image classification models, the system comprising:

a training image source comprising training images associated with labels, wherein the labels include codes used by a trademark registration organization or an industrial design registration organization, and wherein the labels have a hierarchical structure having one or more hierarchy levels;

an image processing device communicably coupled to the training image source, wherein the image processing device is configured to obtain the training images from the training image source and to train a separate deep convolutional neural network for each hierarchy level of the labels using the training images and the labels associated with the training images; and

an input image source comprising input images to be classified into image classification classes, wherein classifying an input image into one or more image classification classes associates the input image with one or more labels corresponding to the one or more image classification classes, wherein probabilities of accurately classifying other images into the one or more classes is pre-computed for use in classifying an input image into the one or more classes.

17. The system of Claim 16, wherein a classification layer of the deep convolutional neural network is based on soft-sigmoid activation.

18. The system of Claim 16, wherein the labels include one or more of U.S. design codes used by the U.S. Patent and Trademark Office, Vienna codes, keywords image constituents, or codes used in Locarno classification.