



- (51) **International Patent Classification:**
G06N 3/08 (2006.01) *G06F 17/00* (2006.01)
- (21) **International Application Number:**
PCT/US2016/027193
- (22) **International Filing Date:**
13 April 2016 (13.04.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/148,788 17 April 2015 (17.04.2015) US
62/251,425 5 November 2015 (05.11.2015) US
15/095,260 11 April 2016 (11.04.2016) US
- (71) **Applicant:** NEC LABORATORIES AMERICA, INC.
[US/US]; 4 Independence Way, Suite 200, Princeton, New Jersey 08540 (US).
- (72) **Inventors:** ZHOU, Feng; 9500 Gilman Drive, LaJolla, California 92093 (US). LIN, Yuanqing; 4 Independence Way, Suite 200, Princeton, New Jersey 08540 (US).
- (74) **Agent:** KOLODKA, Joseph; NEC Laboratories America, Inc., 4 Independence Way, Suite 200, Princeton, New Jersey 08540 (US).

(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

— with international search report (Art. 21(3))

(54) **Title:** FINE-GRAINED IMAGE CLASSIFICATION BY EXPLORING BIPARTITE-GRAPH LABELS




FIG. 1

(57) **Abstract:** Systems and methods are disclosed for deep learning and classifying images of objects by receiving images of objects for training or classification of the objects; producing fine-grained labels of the objects; providing object images to a multi-class convolutional neural network (CNN) having a softmax layer and a final fully connected layer to explicitly model bipartite-graph labels (BGLs); and optimizing the CNN with global back-propagation.

FINE-GRAINED IMAGE CLASSIFICATION BY EXPLORING BIPARTITE-GRAPH LABELS

This application claims priority to Provisional Application 62/251,425 filed 11/5/2015 and 62/148,788 filed 4/17/2015, the content of which is incorporated by reference.

BACKGROUND

The present system relates to fine-grained image classification of real-world physical objects.

Recognition of objects from real-world images is an important task that has many applications, including smart car, security, learning associated with physical structures and products, among others. Fine-grained image classification needs to discern subtle differences among similar classes. The majority of existing approaches have thus been focusing on localizing and describing discriminative object parts in fine-grained domains. Various pose-normalization pooling strategies combined with 2D or 3D geometry have been preferred for recognizing birds. The main drawback of these approaches is that part annotations are significantly more challenging to collect than image labels. Instead, a variety of methods have been developed towards the goal of finding object parts in an unsupervised or semi-supervised fashion.

To provide good features for recognition, another prominent direction is to adopt detection and segmentation methods as an initial step and to filter out noise and clutter in background. However, better feature through segmentation always comes with computational cost as segmentation is often computationally expensive.

While most existing works focus on single-label classification problem, it is more natural to describe real world images with multiple labels like tags or attributes. According to the assumptions on label structures, previous work on structural label learning can be roughly categorized as learning binary, relative or hierarchical attributes.

Much prior work focuses on learning binary attributes that indicate the presence of a certain property in an image or not. For instance, previous works have shown the benefit of learning binary attributes for face verification, texture recognition, clothing searching, and zero-

shot learning. However, binary attributes are restrictive when the description of certain object property is continuous or ambiguous.

To address the limitation of binary attributes, comparing attributes has gained attention in the last years. The relative-attribute approach learns a global linear ranking function for each attribute, offering a semantically richer way to describe and compare objects. While a promising direction, a global ranking of attributes tends to fail when facing fine-grained visual comparisons. One existing system provides learning local functions that tailor the comparisons to neighborhood statistics of the data. A Bayesian strategy can infer when images are indistinguishable for a given attribute.

SUMMARY

In one aspect, systems and methods are disclosed for deep learning and classifying images of objects by receiving images of objects for training or classification of the objects; producing fine-grained labels of the objects; providing object images to a multi-class convolutional neural network (CNN) having a softmax layer and a final fully connected layer to explicitly model bipartite-graph labels (BGLs); and optimizing the CNN with global back-propagation.

Advantages of the system may include one or more of the following. The method achieved state-of-the-arts accuracy in many benchmark dataset. The method is efficient enough to handle very large-scale dataset. The key difference is the use of the bipartite structure of labels. Three benefits follow: (1) it leverages the inherent sub-structure among labels to learn a more accurate model; (2) it regularizes the learning process to alleviate the overfitting problem in CNN; and (3) it is computationally efficient in inference thanks to the bipartite structure. In various tests, the system can predict the instance-level sub-ordinate categories of cars (eg., Honda Accord 2012) and food (Mapo Tofu made by a particular restaurant) from real-world images, which are even difficult to recognize for humans. Focusing on different goal on exploring label structures, our method can be potentially combined with other methods to further improve the performance.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates three ultra-fine grained classes.

FIGS. 2A-2C shows exemplary comparison of output layers, with FIG. 2A showing Softmax (SM), FIG. 2B showing a BGL with one feature (BGL1) and FIG. 2C showing a BGL with multiple features (BGL m).

FIG. 3 shows an exemplary system for detecting food from a particular restaurant source.

FIG. 4 shows exemplary CNNs with the bipartite-graph labels.

Fig. 5A-5F show exemplary comparisons for the preferred BGL system with different baselines.

FIG. 6 shows an exemplary system to run the CNNs with the bipartite-graph labels.

DESCRIPTION

Figure 1. Illustration of three ultra-fine grained classes (middle), Mapo Tofu of Restaurant A, Salt Pepper Tofu of Restaurant B, and Mapo Tofu of Restaurant B. Their relationships can be modeled through three bipartite graphs, fine-grained classes vs. general food dishes (left) and fine-grained classes vs. two ingredients (right). The system incorporates the rich bipartite-graph labels (BGL) into convolutional neural network training to improve recognition accuracy. The preferred BGL method is used in a common multi-class convolutional neural network (CNN) framework, which is compatible to most popular architectures such as AlexNet, GoogLeNet and VGGNet. The approach modifies their softmax layer as well as the last fully connected layer to explicitly model BGLs. The resulted CNN is optimized by a global back-propagation.

In the following, bold capital letters denote a matrix $\mathbf{X}$, bold lower-case letters a column vector $\mathbf{x}$. All non-bold letters represent scalars. $\mathbf{x}_i$ represents the i^{th} column of the matrix $\mathbf{X}$. x_{ij} denotes the scalar in the i^{th} row and j^{th} column of the matrix $\mathbf{X}$. $1_{[i=j]}$ is an indicator function and its value equals to 1 if $i = j$ and 0 otherwise. The system receives a set of n images $X = \{(\mathbf{x}, y), \dots\}$ for training, where each image $\mathbf{x}$ is annotated with one of k fine-grained labels, $y = 1, \dots, k$. In one embodiment, $\mathbf{x} \in \mathbb{R}^d$ denotes the input feature of the last fully-connected layer, which generates k scores $\mathbf{f} \in \mathbb{R}^k$ through a linear function $\mathbf{f} = \mathbf{W}^T \mathbf{x}$ defined by the parameters $\mathbf{W} \in \mathbb{R}^{d \times k}$. In a nutshell, the last layer of CNN is to minimize the negative log-likelihood over the training data,

$$\min_{\mathbf{W}} \sum_{(\mathbf{x}, y) \in X} -\log P(y|\mathbf{x}, \mathbf{W}), \quad (1)$$

where the softmax score,

$$P(i|\mathbf{x}, \mathbf{W}) = \frac{e^{f_i}}{\sum_{j=1}^k e^{f_j}} \doteq p_i, \quad (2)$$

encodes the posterior probability of image being classified as the fine-grained class.

Despite the great improvements achieved on base-class recognition in the last few years, recognizing object classes in ultra-fine granularity like the example shown in Fig. 1 is still challenging mainly for two reasons. First, unlike general recognition task, the amount of labels with ultra-fine granularity is often limited. The training of a high-capacity CNN model is thus more prone to overfitting. Second, it is difficult to learn discriminative features to differentiate between similar fine-grained classes in the presence of large within-class difference.

To address these issues, we use bipartite-graph labels (BGL) to jointly model the fine-grained classes with pre-defined *coarse classes*. Generally speaking, the choices of coarse classes can be any grouping of fine-grained classes. Typical examples include bigger classes, attributes or tags. For instance, Fig. 1 shows three types of coarse classes defined on the fine-grained Tofu dishes (middle). In the first case (Graph 1 in Fig. 1), the coarse classes are two general Tofu food classes by neglecting the restaurant tags. In the second and third cases (Graph 2 and 3 in Fig. 1), the coarse classes are binary attributes according to the presence of some ingredient in the dishes. Compared to the original softmax loss (Eq. 3.1) defined only on fine-grained labels, the introduction of coarse classes in BGL has three benefits: (1) New coarse classes bring in additional supervised information so that the fine-grained classes connected to the same coarse class can share training data with each other. (2) All fine-grained classes are implicitly ranked according to the connection with coarse classes. For instance, Toyota Camry 2014 and Toyota Camry 2015 are much closer to each other compared to Honda Accord 2015. This hierarchical ranking guides CNN training to learn more discriminative features to capture subtle difference between fine-grained classes. (3) Compared to image-level labels (, image attribute, bounding box, segmentation mask) that are expensive to obtain, it is relatively cheaper and more natural to define coarse classes over fine-grained labels. This fact endows BGL the board applicability in real scenario.

Given m types of coarse classes, where each type j contains k_j coarse classes, BGL models their relations with the k fine-grained classes as m bipartite graphs grouped in a star-like structure. Take Fig. 1 for instance, where the three types of coarse classes form three separated bipartite graphs with the fine-grained Tofu dishes, and there is no direct connection among the three types of coarse classes. For each graph of coarse type j , we encode its

bipartite structure in a binary association matrix $\mathbf{G}_j \in \{0,1\}^{k \times k_j}$, whose element $g_{ic_j}^j = 1$ if the i^{th} fine-grained label is connected with coarse label c_j . As it will become clear later, this star-like composition of bipartite graphs enables BGL to perform exact inference as opposed to the use of other general label graphs.

To generate the scores $\mathbf{f}_j = \mathbf{W}_j^T \mathbf{x} \in \mathbb{R}^{k_j}$ for coarse classes of type j , we augment the last fully-connected layer with m additional variables, $\{\mathbf{W}_j\}_j$, where $\mathbf{W}_j \in \mathbb{R}^{d \times k_j}$. Given an input image $\mathbf{x}$ of i^{th} fine-gained class, BGL models its joint probability with any m coarse labels $\{c_j\}_j$ as,

$$P(i, \{c_j\}_j | \mathbf{x}, \mathbf{W}, \{\mathbf{W}_j\}_j) = \frac{1}{z} e^{f_i} \prod_{j=1}^m g_{ic_j}^j e^{f_{c_j}^j},$$

where z is the partition function computed as:

$$z = \sum_{i=1}^k e^{f_i} \prod_{j=1}^m \sum_{c_j=1}^{k_j} g_{ic_j}^j e^{f_{c_j}^j}.$$

At first glance, computing z is infeasible in practice. Because of the bipartite structure of the label graph, however, we could denote the non-zero element in i^{th} row of $\mathbf{G}_j$ as $\phi_i^j = c_j$ where $g_{ic_j}^j = 1$. With this auxiliary function, the computation of z can be simplified as

$$z = \sum_{i=1}^k e^{f_i} \prod_{j=1}^m e^{f_{\phi_i^j}^j}. \quad (3)$$

Given z , the marginal posterior probability over fine-grained and coarse labels can be computed as:

$$P(i | \mathbf{x}, \mathbf{W}, \{\mathbf{W}_j\}_j) = \frac{1}{z} e^{f_i} \prod_{j=1}^m e^{f_{\phi_i^j}^j} \doteq p_i,$$

$$P(c_j | \mathbf{x}, \mathbf{W}, \{\mathbf{W}_l\}_l) = \frac{1}{z} \sum_{i=1}^k g_{ic_j}^j e^{f_i} \prod_{l=1}^m e^{f_{\phi_i^l}^l} \doteq p_{c_j}^j.$$

As discussed before, one of the difficulties in training CNN is the possibility of being overfitting. One common solution is to add a l_2 weight decay term, which is equivalent to sample the columns of $\mathbf{W}$ from a Gaussian prior. Given the connection among fine-grained and coarse classes, BGL provides another natural hierarchical prior for sampling the weights:

$$P(\mathbf{W}, \{\mathbf{W}_j\}_j) = \prod_{i=1}^k \prod_{j=1}^m \prod_{c_j=1}^{k_j} e^{-\frac{\lambda}{2} g_{i,c_j}^j \|\mathbf{w}_i - \mathbf{w}_{c_j}^j\|^2} \triangleq p_w.$$

This prior expects $\mathbf{w}_i$ and $\mathbf{w}_{c_j}^j$ have small distance if i^{th} fine-grained label is connected to coarse class c_j of type j . BGL can handle general coarse labels such as multiple attributes.

In summary, given the training data $\mathbf{X}$ and the graph label defined by $\{\mathbf{G}_j\}_j$, the last layer of CNN with BGL aims to minimize the joint negative log-likelihood with proper regularization over the weights:

$$\min_{\mathbf{W}, \{\mathbf{W}_j\}_j} \sum_{(\mathbf{x}, y) \in \mathcal{X}} \left(-\log p_y - \sum_{j=1}^m \log p_{\phi_q^j}^j \right) - \log p_w. \quad (4)$$

Figure 2a-2c shows various exemplary comparison of output layers: (a) Softmax (SM), (b) BGL with one feature (BGL), and (c) BGL with multiple features (BGL). In one implementation, we optimized BGL using the standard back-propagation with mini-batch stochastic gradient descent. The gradients for each parameter can be computed all in closed-form:

$$\begin{aligned}
\frac{\partial \log p_y}{\partial f_i} &= 1_{[i=y]} - p_i, \quad \frac{\partial \log p_y}{\partial f_{c_j}^j} = 1_{[g_{yc_j}^j=1]} - p_{c_j}^j, \\
\frac{\partial \log p_{\phi_y^j}^j}{\partial f_i} &= \frac{p_i}{p_{\phi_y^j}^j} 1_{[g_{iy}^j=1]} - p_i, \quad \frac{\partial \log p_{\phi_y^j}^j}{\partial f_{c_j}^j} = 1_{[c_j=\phi_y^j]} - p_{c_j}^j, \\
\frac{\partial \log p_{\phi_y^j}^j}{\partial f_{c_l}^l} &= \sum_{i=1}^k g_{i\phi_y^j}^l g_{il}^l \frac{p_i}{p_{\phi_y^j}^j} - p_{c_l}^l, l \neq j,
\end{aligned} \tag{5}$$

$$\frac{\partial \log p_w}{\partial \mathbf{w}_i} = -\lambda \sum_{j=1}^m \sum_{c_j=1}^{b_j} g_{ic_j}^j (\mathbf{w}_i - \mathbf{w}_{c_j}^j), \tag{6}$$

$$\frac{\partial \log p_w}{\partial \mathbf{w}_{c_j}^j} = -\lambda \sum_{i=1}^k g_{ic_j}^j (\mathbf{w}_{c_j}^j - \mathbf{w}_i). \tag{7}$$

The system implements the following in one implementation: (1) Computing $p_i/p_{\phi_y^j}^j$ by independently calculating p_i and $p_{\phi_y^j}^j$ is not numerically stable because $p_{\phi_y^j}^j$ could be very small. It is better to jointly normalize the two terms first. (2) Directly computing Eq. 3.3 has a quadratic complexity with respect to the number of coarse classes. But it can be reduced to linear because most computations are redundant. (3) we assume the same feature $\mathbf{x}$ is used for computing both the fine-grained $\mathbf{f} = \mathbf{W}^T \mathbf{x}$ and coarse scores $\mathbf{f}_j = \mathbf{W}_j^T \mathbf{x}$. In fact, BGL can naturally combine multiple CNNs as shown in Fig. 2c. This allows the model to learn different low-level features $\mathbf{x}_j$ for coarse labels $\mathbf{f}_j = \mathbf{W}_j^T \mathbf{x}_j$.

FIG. 3 shows an exemplary system for detecting food from a particular restaurant source. The system is well suited for deep learning and classifying images of objects by receiving images of objects for training or classification of the objects; producing fine-grained labels of the objects; providing object images to a multi-class convolutional neural network (CNN) having a softmax layer and a last fully connected layer to explicitly model bipartite-graph labels (BGLs); and optimizing the CNN with global back-propagation. In one embodiment, the system includes generating bipartite-graph labels for fine-grained object recognition in the CNNs and modeling a bi-partite graph structure in a conditional random fields (CRF) and optimizing low-level features and high-level connections in the CNN.

Fig. 4 shows one implementation of BGL with the off-the-shelf Caffe platform. The system is then tested on three popular CNN frameworks, AlexNet (**AN**), GoogLeNet (**GN**), and VGGNet (**VGG**). For each of them, we compared three settings: **SM** with the traditional softmax loss; **BGL₁** by modifying the softmax layer and the last fully connected layer as the preferred BGL approach; and **BGL_m** by combining multiple networks that generate different features for fine-grained and coarse classes.

In sum, the preferred BGL exploits the rich class relationships in the very challenging ultra-fine grained tasks. BGL improves the traditional softmax loss by jointly modeling fine-grained and coarse labels through bipartite-graph labels. The use of a special bipartite structure enables BGL to be efficient in inference. We also contribute Food-975, an ultra-fine grained food recognition benchmark dataset. We show that the preferred BGL approach improves over previous work on a variety of datasets.

Our models were trained for 100 epochs on a single NVIDIA K40 GPU. We adopted the default hyper-parameters as used by Caffe. In all experiments, we fine-tuned from pre-trained ImageNet model. During training, we down-sampled the images to a fixed 256-by-256 resolution, from which we randomly cropped 224-by-224 patches for training. We also did their horizontal reflection for further data augmentation. During testing, we evaluated the top-1 accuracy using two cropping strategies: (1) single-view (**SV**) by cropping the center 224-by-224 patch of the testing image, and (2) multi-view (**MV**) by averaging the center, 4 corners and their mirrored versions. In the first three datasets, we evaluated our methods using two protocols, without (**w/o. BBox**) and with (**w/. BBox**) the use of ground-truth bounding box to crop out the object both at training and testing.

The first experiment validates our approach on the Stanford car dataset, which contains 16,185 images of 196 car categories. We adopted a 50-50 split by dividing the data into 8,144 images for training and the rest for testing. Each category is typically at the level of maker, model and year, , Audi A5 Coupe 12. We manually assigned each fine-grained label to one of 9 coarse body types. Our method of VGG-BGL_m beats all previous works except one which leveraged the part information in an unsupervised way. BGL has the advantage of predicting coarse label. For instance, GN-BGL_m achieved 95.7% in predicting the type of a car image.

In the second experiment, we tested our method on CUB-200-2011, which is generally considered the most competitive dataset within fine-grained recognition. CUB-200-2011 contains 11,788 images of 200 bird species. We used the provided train/test split and reported results in terms of classification accuracy. To get the label hierarchy, we adopted the annotated 312 visual attributes associated with the dataset. These attributes are divided into 28 groups, each of which has about 10 choices. According to the provided confidence score, we assigned each attribute group with the most confident choice for each bird specie. From the experiments, we observed using BGL modules can consistently improved AN-SM, GN-SM and VGG-SM. Without any pose alignment steps, our method GN-BGL_m obtained 76.9% without the use of bounding box, improving the recent part-based method by 3%. In addition, GN-BGL_m achieved 89.3% and 93.3% accuracy on predicting attributes of wing color and shape. However, our method still performed worse than the latest methods which show the significant advantage by exploring part information for bird recognition. BGL improves the last fully connected and loss layer for attribute learning, while the other systems focus on integrating object part information into convolutional layers. Therefore, it is possible to combine these two orthogonal efforts to further improve the overall performance.

In the third experiment, we test our method on the Car-333 dataset, which contains 157,023 training images and 7,840 testing images. Compared to the Stanford car dataset, the images in Car-333 were end-user photos and thus more naturally photographed. Each of the 333 labels is composed by maker, model and year range. Two cars of the same model but manufactured in different year ranges are considered different classes. To test BGL, we generated two sets of coarse labels: 10 "type" coarse labels manually defined according to the geometric shape of each car model and 140 model coarse labels by aggregating year range labels. The bounding box of each image was generated by Regionlets, the state-of-the-art object detection method.

We tested BGL with three combinations of the coarse labels: using either model or type, and using model and type jointly. In particular, BGL gains much more improvements using the 140 model coarse labels than the 10 type labels. This is because the images of the cars of the same "model" are more similar than the ones in the same "type" and it defines richer

relationships among fine-grained classes. Nevertheless, BGL can still get benefit from putting the ``type'' labels on top of the ``model'' labels to form a three-level label hierarchy. Finally, GN-BGL_m significantly improved the performance of GN-SM from 79.8% to 86.4% without the use of bounding box. For more result on AN and VGG, please refer to the supplementary material.

Compared to GN-SM, GN-BGL needs only very little additional computation to perform exact inference in the loss function layer. This demonstrates the efficiency of modeling label dependency in a bipartite graphs. For the last fully connected (FC) layer, BGL performs exactly the same computation as GN in the forward passing, but needs additional cost for updating the gradient (Eq. 3.3 and Eq. 3.3) in the backward passing. Because both the loss and last FC layers take a very small portion of the whole pipeline, we found the total time difference between BGL and GN was minor.

Figs. 5A-5F shows an exemplary performance illustration on the Food-975 dataset. Fig. 5A shows exemplary distribution of training images for each fine-grained (top), dish (middle), and ingredient (bottom) category. Fig. 5B shows exemplary testing image. Fig. 5C shows exemplary ingredients predicted by GN-BGL_m. Fig. 5D shows exemplary top-5 predictions generated by GN-SM (top) and the preferred GN-BGL_m (bottom) respectively. Fig. 5E shows similar training exemplars according to the input feature $\mathbf{x}$. Fig. 5F illustrates the accuracy of the food recognition. Given a food image from a restaurant, the system can answer the question “which restaurant which dish”? This is an ultra-fine grained image recognition problem and is very challenging. Such ultra-fine granularity brings along rich relationships among the classes. In particular, different restaurants may cook similar dishes with different ingredients.

We collected a high quality food dataset. We sent 6 data collectors to 6 restaurants. Each data collector was in charge of one restaurant and took the photo of almost every dish the restaurant had cooked during a period of 1-2 months. Finally, we captured 32135 high-resolution food photos of 975 menu items from the 6 restaurants for training. We evaluated our method in two settings. To test in a *controlled* setting, we took additional 4951 photos in different days. To mimic a realistic scenario in the *wild*, we downloaded 351 images from yelp.com posted by consumers visiting the same restaurants. To model the class relationship,

we created a three-level hierarchy. In the first level, we have the 975 fine-grained labels; in the middle, we created 781 different dishes by aggregating restaurant tags; at last, we came up a detailed list of 51 ingredient attributes that precisely describes the food composition.

Given a food image, can a fine-grained object recognition engine tell "which restaurant which dish" the food belongs to? Such ultra-fine grained image recognition is the key for many applications like search by images, but it is very challenging because it needs to discern subtle difference between classes while dealing with the scarcity of training data. Fortunately, the ultra-fine granularity naturally brings rich relationships among object classes. The system exploits the rich relationships through bipartite-graph labels (BGL). BGL is modeled in an overall convolutional neural networks and the resulting system can be optimized through back-propagation. The system is computationally efficient in inference thanks to the bipartite structure. To facilitate the study, we construct a new food benchmark dataset, which consists of 37,885 food images collected from 6 restaurants and totally 975 menus. Experimental results on this new food and three other datasets demonstrates BGL advances previous works in fine-grained object recognition. Fig. 5A-5F compared the preferred BGL approach with different baselines. Fig. 5E compares our method with AN-SM, GN-SM and VGG-SM in both the controlled and wild settings. BGL approach consistently out-performed AN and GN in both settings. This indicates the effectiveness of exploring the label dependency in ultra-fine grained food recognition. Interestingly, by using both the dish and ingredient labels, BGL can gain a much larger improvement than only using one of them. This implies the connections between the dish labels and the ingredient labels have very useful information. Overall, by exploring the label dependency, the preferred BGL approach consistently achieved 6-8% improvement from GN baseline at different training setups.

For the ultra-fine grained image recognition, the system can handle many more classes. For example, we are constructing a large food dataset from thousands of restaurants where the number of ultra-fine grained food classes can grow into hundreds of thousands. We believe that the research in this direction, ultra-fine-grained image recognition (recognizing images almost on instance-level), holds the key for using images as a media to retrieve information, which is often called *search by image*. The label structure can be manually defined or can be

learned during training. While discrete labels in BGL are discussed, the system can work with continuous labels (for example, regression or ranking). Instead of operating only at class level, we can generalize BGL to deal with image-level labels. This can make the performance of BGL more robust in the case when the attribute label is ambiguous for fine-grained class.

FIG. 6 with an exemplary processing system 100, to which the present principles may be applied, is illustratively depicted in accordance with an embodiment of the present principles, for deep learning and classifying images of objects by receiving images of objects for training or classification of the objects; producing fine-grained labels of the objects; providing object images to a multi-class convolutional neural network (CNN) having a softmax layer and a final fully connected layer to explicitly model bipartite-graph labels (BGLs); and optimizing the CNN with global back-propagation.

The processing system 100 includes at least one processor (CPU) 104 operatively coupled to other components via a system bus 102. A cache 106, a Read Only Memory (ROM) 108, a Random Access Memory (RAM) 110, an input/output (I/O) adapter 120, a sound adapter 130, a network adapter 140, a user interface adapter 150, and a display adapter 160, are operatively coupled to the system bus 102.

A first storage device 122 and a second storage device 124 are operatively coupled to system bus 102 by the I/O adapter 120. The storage devices 122 and 124 can be any of a disk storage device (e.g., a magnetic or optical disk storage device), a solid state magnetic device, and so forth. The storage devices 122 and 124 can be the same type of storage device or different types of storage devices.

A speaker 132 is operatively coupled to system bus 102 by the sound adapter 130. A transceiver 142 is operatively coupled to system bus 102 by network adapter 140. A display device 162 is operatively coupled to system bus 102 by display adapter 160.

A first user input device 152, a second user input device 154, and a third user input device 156 are operatively coupled to system bus 102 by user interface adapter 150. The user input devices 152, 154, and 156 can be any of a keyboard, a mouse, a keypad, an image capture device, a motion sensing device, a microphone, a device incorporating the functionality of at least two of the preceding devices, and so forth. Of course, other types of input devices can

also be used, while maintaining the spirit of the present principles. The user input devices 152, 154, and 156 can be the same type of user input device or different types of user input devices. The user input devices 152, 154, and 156 are used to input and output information to and from system 100.

Of course, the processing system 100 may also include other elements (not shown), as readily contemplated by one of skill in the art, as well as omit certain elements. For example, various other input devices and/or output devices can be included in processing system 100, depending upon the particular implementation of the same, as readily understood by one of ordinary skill in the art. For example, various types of wireless and/or wired input and/or output devices can be used. Moreover, additional processors, controllers, memories, and so forth, in various configurations can also be utilized as readily appreciated by one of ordinary skill in the art. These and other variations of the processing system 100 are readily contemplated by one of ordinary skill in the art given the teachings of the present principles provided herein.

While the machine-readable storage medium is shown in an exemplary embodiment to be a single medium, the term “machine-readable storage medium” should be taken to include a single medium or multiple media (e.g., a centralized or distributed database, and/or associated caches and servers) that store the one or more sets of instructions. The term “machine-readable storage medium” shall also be taken to include any medium that is capable of storing or encoding a set of instructions for execution by the machine and that cause the machine to perform any one or more of the methodologies of the present invention. The term “machine-readable storage medium” shall accordingly be taken to include, but not be limited to, solid-state memories, and optical and magnetic media.

It is to be understood that the above description is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reading and understanding the above description. Although the present invention has been described with reference to specific exemplary embodiments, it will be recognized that the invention is not limited to the embodiments described, but can be practiced with modification and alteration within the spirit and scope of the appended claims. Accordingly, the specification and

drawings are to be regarded in an illustrative sense rather than a restrictive sense. The scope of the invention should, therefore, be determined with reference to the appended claims, along

What is claimed is:

1. A method for deep learning and classifying images of objects, comprising:
 receiving images of objects for training or classification of the objects;
 producing fine-grained labels of the objects;
 providing object images to a multi-class convolutional neural network (CNN) having a softmax layer and a final fully connected layer to explicitly model bipartite-graph labels (BGLs); and
 optimizing the CNN with global back-propagation.
2. The method of claim 1, comprising modeling a relation between the fine-grained labels and attributes in a hierarchical manner.
3. The method of claim 1, comprising applying the BGLs to jointly model fine-grained classes with predefined coarse classes.
4. The method of claim 3, comprising grouping fine-grained classes as a coarse class.
5. The method of claim 1, comprising m types of coarse classes, where each type j contains k_j coarse classes, BGL models their relations with the k fine-grained classes as m bipartite graphs grouped in a star-like structure.
6. The method of claim 1, for each graph of coarse type j, comprising encoding a bipartite structure in a binary association matrix $\mathbf{G}_j \in \{0, 1\}^{k \times k_j}$, whose element $g_{ic_j}^j = 1$ if the i-th fine-grained label is connected with coarse label c_j .
7. The method of claim 1, comprising:
 generating scores $\mathbf{f} = \mathbf{W}^T \mathbf{x} \in \mathbb{R}^k$ for k fine-grained classes;

generating scores $\mathbf{f}_j = \mathbf{W}_j^T \mathbf{x} \in \mathbb{R}^{k_j}$ for k_j coarse classes of type j ;
 augmenting a last fully-connected layer $\mathbf{W}$ with m additional variables, $\{\mathbf{W}_j\}_{j=1}^m$ and
 with an input image $\mathbf{x}$ of i -th fine-gained class, modeling a BGL joint probability with m
 coarse labels $\{c_j\}_{j=1}^m$ as,

$$P(i, \{c_j\}_j | \mathbf{x}, \mathbf{W}, \{\mathbf{W}_j\}_j) = \frac{1}{z} e^{f_i} \prod_{j=1}^m g_{ic_j}^j e^{f_{c_j}^j}$$

where the z is the partition function computed as,

$$z = \sum_{i=1}^k e^{f_i} \prod_{j=1}^m \sum_{c_j=1}^{k_j} g_{ic_j}^j e^{f_{c_j}^j}.$$

8. The method of claim 1, comprising optimizing the BGL using back-propagation with mini-batch stochastic gradient descent.

9. The method of claim 8, comprising determining gradients for each parameter in closed-form.

10. The method of claim 9, comprising determining the gradients as:

$$\begin{aligned} \frac{\partial \log p_y}{\partial f_i} &= 1_{[i=y]} - p_i, \quad \frac{\partial \log p_y}{\partial f_{c_j}^j} = 1_{[g_{yc_j}^j=1]} - p_{c_j}^j, \\ \frac{\partial \log p_{\phi_y^j}^j}{\partial f_i} &= \frac{p_i}{p_{\phi_y^j}^j} 1_{[g_{i\phi_y^j}^j=1]} - p_i, \quad \frac{\partial \log p_{\phi_y^j}^j}{\partial f_{c_j}^j} = 1_{[c_j=\phi_y^j]} - p_{c_j}^j, \\ \frac{\partial \log p_{\phi_y^j}^j}{\partial f_{c_l}^l} &= \sum_{i=1}^k g_{i\phi_y^j}^j g_{il}^l \frac{p_i}{p_{\phi_y^j}^j} - p_{c_l}^l, l \neq j, \\ \frac{\partial \log p_w}{\partial \mathbf{w}_i} &= -\lambda \sum_{j=1}^m \sum_{c_j=1}^{k_j} g_{ic_j}^j (\mathbf{w}_i - \mathbf{w}_{c_j}^j), \\ \frac{\partial \log p_w}{\partial \mathbf{w}_{c_j}^j} &= -\lambda \sum_{i=1}^k g_{ic_j}^j (\mathbf{w}_{c_j}^j - \mathbf{w}_i). \end{aligned}$$

where $\phi_y^j = c_y$ denotes the non-zero element of i -th row of $\mathbf{G}_j$, that is, $g_{ic_j}^j = 1$; p_i and $p_{\phi_y^j}^j$ are the marginal probability derived from claim 7; $\mathbf{w}_i$ is i -th column of $\mathbf{W}$; $\mathbf{w}_{c_j}^j$ is c_j -column of $\mathbf{W}_j$.

11. The method of claim 1, wherein the objects comprises a finished good, further comprising modeling label structures between fine-grained labels and their ingredients in a bipartite graph.

12. The method of claim 1, wherein the finished good comprises food.

13. The method of claim 1, comprising classifying food type and restaurant source.

14. The method of claim 1, comprising modeling a bi-partite graph structure in a conditional random field (CRF) and optimizing lower-level features and high-level connections in the CNN.

15. The method of claim 1, for a training data $\mathbf{X}$ and a graph label defined by $\{\mathbf{G}_j\}_{j=1}^m$, comprising determining a last layer of the CNN with BGL to minimize a joint negative log-likelihood with proper regularization over weights:

$$\min_{\mathbf{W}, \{\mathbf{W}_j\}_j} \sum_{(x,y) \in \mathcal{X}} \left(-\log p_y - \sum_{j=1}^m \log p_{\phi_y^j}^j \right) - \log p_w.$$

see claim 10 for the definition of each component.

16. The method of claim 1, comprising generating bipartite-graph labels for fine-grained object recognition in the CNNs and modeling a bi-partite graph structure in a conditional random fields (CRF) and optimizing low-level features and high-level connections in the CNN.

17. A system to perform deep learning and classifying images of objects, comprising:

a processor;

an image sensor coupled to the processor, the image sensor capturing images of objects;

a multi-class convolutional neural network (CNN) having a softmax layer and a final fully connected layer to explicitly model bipartite-graph labels (BGLs); and

a computer readable module coupled to the CNN for:

receiving images of objects for training or classification of the objects;

producing fine-grained labels of the objects;

providing object images to the multi-class convolutional neural network (CNN) having a softmax layer and a final fully connected layer to explicitly model bipartite-graph labels (BGLs); and

optimizing the CNN with global back-propagation.

18. The system of claim 17, comprising code for modeling a relation between the fine-grained labels and attributes in a hierarchical manner.

19. The system of claim 17, comprising code for applying the BGLs to jointly model fine-grained classes with predefined coarse classes.

20. The system of claim 17, comprising code for:

generating scores $\mathbf{f} = \mathbf{W}^T \mathbf{x} \in \mathbb{R}^k$ for k fine-grained classes;

generating scores $\mathbf{f}_j = \mathbf{W}_j^T \mathbf{x} \in \mathbb{R}^{k_j}$ for k_j coarse classes of type j ;

augmenting a last fully-connected layer $\mathbf{W}$ with m additional variables, $\{\mathbf{W}_j\}_{j=1}^m$ and

with an input image $\mathbf{x}$ of i -th fine-gained class, modeling a BGL joint probability with m coarse labels $\{c_j\}_{j=1}^m$ as,

$$P(i, \{c_j\}_j | \mathbf{x}, \mathbf{W}, \{\mathbf{W}_j\}_j) = \frac{1}{Z} e^{f_i} \prod_{j=1}^m g_{ic_j}^j e^{f_{c_j}^j}$$




FIG. 1




FIG. 2A

FIG. 2B

FIG. 2C




FIG. 3




FIG. 4




FIG. 5A

A: Fish Stew Preserved Veg.
└ Fish Stew Preserved Veg.
└ Cilantro └ Chili └ Fish └ Veg.




FIG. 5B

F: Steamed Clams Egg
A: Fish Stew Preserved Veg.
A: Chongqing Mao Xie Wang
A: Fish Fillet In Szechuan
C: Duo Zui Zheng Yu Tou Nan


GN-SM

A: Fish Stew Preserved Veg.
F: Steamed Clams Egg
A: Chongqing Mao Xie Wang
A: Fish Fillet In Szechuan
F: Crispy Fish Fillet Sea Moss

GN-BGL_m

FIG. 5C

F: Steamed Clams Egg
└ Steamed Clams Egg
└ Egg └ Green onion └ Seafood



A: Fish Stew Preserved Veg.
└ Fish Stew Preserved Veg.
└ Cilantro └ Chili └ Fish └ Veg.



FIG. 5D

FIG. 5E

Accuracy(%)	Controlled		Wild	
	SV	MV	SV	MV
AN-SM	84.1	84.7	47.9	49.0
GN-SM	87.2	87.5	56.1	55.6
VGG-SM	87.3	87.7	51.0	54.4
Dish Coarse Labels				
GN-BGL ₁	87.5	88.0	55.8	57.8
GN-BGL _m	87.8	88.0	58.4	59.0
Ingredient Coarse Labels				
GN-BGL ₁	87.9	88.0	58.0	57.3
GN-BGL _m	88.4	88.6	59.0	57.3
Both Coarse Labels				
GN-BGL ₁	88.0	88.4	62.1	61.8
GN-BGL _m	89.0	88.7	62.3	62.0

FIG. 5F




FIG.6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/027193**A. CLASSIFICATION OF SUBJECT MATTER****G06N 3/08(2006.01)i, G06F 17/00(2006.01)i**

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)

G06N 3/08; G06F 17/30; G06K 9/00; G06K 9/62; G06F 15/18; G06F 17/00

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & keywords: deep learning, classification, fine-grained label, CNN(Convolutional Neural Network), softmax layer, bipartite-graph label, and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	JIANG WANG et al., `Learning Fine-grained Image Similarity with Deep Ranking`, In: 2014 IEEE Conference on Computer Vision and Pattern Recognition, 23 June 2014 <URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6909576&tag=1 > See pages 1-8.	1-20
A	US 2009-0297016 A1 (RICHARD LEVENSON et al.) 03 December 2009 See paragraphs [0064]-[0104]; claim 1; and figure 2.	1-20
A	US 2013-0290222 A1 (ALBERT GORDO et al.) 31 October 2013 See paragraphs [0028]-[0113]; claim 1; and figure 2.	1-20
A	US 2014-0270431 A1 (SONY CORPORATION) 18 September 2014 See paragraph [0020]; claim 1; and figure 1	1-20
A	US 2013-0266214 A1 (BRIGHAM YOUNG UNIVERSITY) 10 October 2013 See paragraphs [0043]-[0049]; claim 1; and figure 5.	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

20 July 2016 (20.07.2016)

Date of mailing of the international search report

20 July 2016 (20.07.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

BYUN, Sung Cheal

Telephone No. +82-42-481-8262



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/027193

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2009-0297016 A1	03/12/2009	CN 101151623 A CN 101916359 A EP 1846869 A1 EP 2237189 A2 US 2006-0245631 A1 US 2011-0255753 A1 US 2013-0077875 A1 US 2014-0376824 A1 US 7555155 B2 US 7953264 B2 US 8280140 B2 US 8639043 B2 WO 2006-081547 A1	26/03/2008 15/12/2010 24/10/2007 06/10/2010 02/11/2006 20/10/2011 28/03/2013 25/12/2014 30/06/2009 31/05/2011 02/10/2012 28/01/2014 03/08/2006
US 2013-0290222 A1	31/10/2013	FR 2990035 A1 US 9075824 B2	01/11/2013 07/07/2015
US 2014-0270431 A1	18/09/2014	WO 2014-144103 A1	18/09/2014
US 2013-0266214 A1	10/10/2013	US 9317779 B2	19/04/2016