



Universiteit
Leiden
The Netherlands

Deep learning for visual understanding

Guo, Y.; Guo Y.

Citation

Guo, Y. (2017, October 5). *Deep learning for visual understanding*. Retrieved from <https://hdl.handle.net/1887/52990>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/52990>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/52990> holds various files of this Leiden University dissertation.

Author: Guo, Y.

Title: Deep learning for visual understanding

Issue Date: 2017-10-05

Chapter 5

CNN-RNN: A Large-scale Hierarchical Image Classification Framework

Objects are often organized in a semantic hierarchy of categories, where fine-level categories are grouped into coarse-level categories according to their semantic relations. While previous works usually only classify objects into the leaf categories, we argue that generating hierarchical labels can actually describe how the leaf categories evolved from higher level coarse-grained categories, thus can provide a better understanding of the objects. In this chapter, we propose to utilize the CNN-RNN framework to address the hierarchical image classification task. CNN allows us to obtain discriminative features for the input images, and RNN enables us to jointly optimize the classification of coarse and fine labels. This framework can not only generate hierarchical labels for images, but also improve the traditional leaf-level classification performance due to incorporating the hierarchical information. Moreover, this framework can be built on top of any CNN architecture which is primarily designed for leaf-level classification. Accordingly, we build a high performance network based on the CNN-RNN paradigm which outperforms the original CNN (wider-ResNet) and also the current state-of-the-art. In addition, we investigate how to utilize the CNN-RNN framework

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

to improve the fine category classification when a fraction of the training data is only annotated with coarse labels. Experimental results demonstrate that CNN-RNN can use the coarse-labeled training data to improve the classification of fine categories, and in some cases it even surpasses the performance achieved by fully annotated training data. This reveals that, CNN-RNN can alleviate the challenge of specialized and expensive annotation of fine labels.

5.1 Introduction

Image classification has long been an active area of research, which aims to classify images into pre-defined categories, and helps people to know what kind of objects the images contain. Traditionally, image classification is mainly performed on small datasets, by encoding local hand-crafted features and using them as input for classifiers [31, 122, 135].

In recent years, two fundamental changes occurred for this task: first, the number of digital images has been increasing exponentially. This brings people more alternatives, and more difficulties, in finding relevant images from this large volume of data. To help people access data, in an effortless and meaningful way, we need a good semantic organization of the categories. Second, deep learning methods have proven to be successful for image classification. In recent years, researchers have built various deep structures [24, 25, 227], and have achieved quite accurate predictions on small datasets [236, 259].

As a consequence, the current research focus has moved to larger and more challenging datasets [260, 261], such as ImageNet [149]. Such datasets often organize the large number of categories in a hierarchy, according to their semantic belongings. The deeper one goes in the hierarchy, the more specific the category is. In contrast to current approaches, which only focus on the leaf categories, we argue that generating hierarchical labels in a coarse-to-fine pattern can present how the semantic categories evolve, and thus can better describe what the objects are. For example, for Figure 5.1(c), the predicted leaf-category label is ‘Triceratops’. Without specialized knowledge, we cannot learn that this category label belongs to the higher level category label ‘Horned Dinosaur’.

The **first contribution** of this work is a framework capable of generating hierarchical labels, by integrating the powerful Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN). CNN is used to generate discriminative features, and RNN is used to generate sequential labels.

There are several notable advantages for the CNN-RNN framework:

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

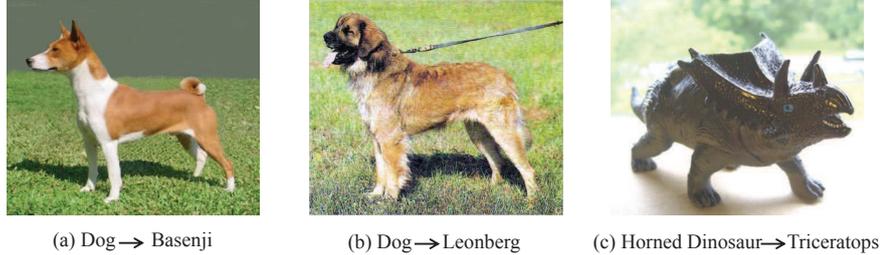


Figure 5.1: The example images with the ‘coarse \rightarrow fine’ label.

(1) Learning things in a hierarchical way is consistent with human perception and concept organization. By predicting the labels in a coarse-to-fine pattern, we can better understand what the objects are, such as depicted in Figure 5.1(c).

(2) It can exploit the relationship between the hierarchical categories, which, in turn, helps the traditional image classification task. For example, when we build the CNN-RNN framework with wrn-28-10 [262], we can increase the accuracy of coarse and fine categories by 2.8% and 1.68%, respectively. To the best of our knowledge, this is the first work trying to employ RNN to improve the classification performance by exploiting the relationship between hierarchical labels.

(3) It is transferrable. In principle, the framework can be built on top of any CNN architecture which is primarily intended for single-level classification, and boost the performance for each hierarchical level. To verify this, we have conducted extensive experiments with three high-performance networks, i.e. CNN-7 [263], wrn-28-10 [262] and our proposed wider-Resnet.

(4) The structure can be trained end-to-end. In contrast to other methods which can only model the category relationship with pre-computed image features [264, 265], the CNN-RNN framework can jointly learn the features and relationship in an end-to-end way, which can improve the final predictions considerably. For a subset of ImageNet 2010, we compared employing pre-computed CNN features to train the RNN with end-to-end training the CNN and RNN and demonstrated a significant improvement of the subcategory accuracy from 77.27% to 82%.

(5) The number of the hierarchical labels can be variable. The flexibility of RNN allows us to generate hierarchical labels of different lengths, i.e. more

specific categories would have more hierarchical labels. We have demonstrated this property on the widely used ImageNet 2012 dataset [21].

As the framework is transferrable, we intend to build a high performance CNN model, and utilize its CNN-RNN variant to further boost the accuracy.

Recently, deep residual networks (ResNet) [227] have attracted great attention because of its leading performance in several image classification tasks and Zagoruyko et al. [262] further presented a thorough experimental study about several important aspects of ResNet, such as the width and depth, and proposed a wide Resnet that obtained better performance than the original ResNet. The **second contribution** of this work is, we build a wider ResNet compared to [262]. Our implementation shows that, the wider-Resnet performs better than [262] on CIFAR-100, and also outperforms the original ResNet with thousands layers. In addition, by utilizing the CNN-RNN framework, we obtain considerably better results than the state-of-the-art.

The performance of deep models has benefited from the accurate and large-scale annotations, such as ImageNet [149]. However, manual labeling is an excessively tedious and expensive task, especially for the fine-grained classes, which often require expert knowledge (e.g. breeds of dogs, flower species, etc.). For example, for Figure 5.1 (a) and Figure 5.1 (b), it is easy to annotate the images with the coarse label ‘dog’, but it requires specialized knowledge to divide them into subcategories ‘Bsenji’ and ‘Leonberg’. One optional thought is, if a part of the training data is only annotated with coarse category labels, whether we could utilize the coarse-labeled training data to improve the classification performance of fine categories?

The **third contribution** of this work is, we investigate how to utilize the CNN-RNN framework to improve the subcategory classification when a fraction of the training data only has coarse labels. By training the CNN-RNN framework on the fully annotated data in the training set, we can exploit the relationship between the coarse and fine categories. Thereby, we can predict the fine labels of the coarse-labeled training data, and then re-train the CNN-RNN model. Experimental results demonstrate that the coarse-labeled training data can normally

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

help the subcategory classification. In some cases, it can even surpass the performance of fully annotated training data. This alleviates the expensive process of fine-grained labeling.

5.2 Related Work

5.2.1 Usage of CNN-RNN framework

Deep learning methods have attracted significant attention [226] and achieved revolutionary successes in various applications [213, 227]. Two important structures for deep learning are CNN and RNN. CNN has proven to be successful in processing image-like data, while RNN is more appropriate in modeling sequential data. Recently, several works [266–271] have attempted to combine them together, and built various CNN-RNN frameworks. Generally, the combination can be divided in two types: the unified combination and the cascaded combination.

The unified combination often attempts to introduce a recurrent property into the traditional CNN structure in order to increase the classification performance. For example, Zuo et al. [269] converted each image into 1D spatial sequences by concatenating the CNN features of different regions, and utilized RNN to learn the spatial dependencies of image regions. Similar work appeared in [272]. The proposed ReNet replaced the ubiquitous convolutional+pooling layer with four recurrent neural networks that sweep horizontally and vertically in both directions across the image. In order to improve the multi-label classification, Wang et al. [271] presented the CNN-RNN framework to learn a joint embedding space in modeling semantic label dependency as well as the image-label relevance.

On the other hand, the cascaded combination would process the CNN and RNN separately, where the RNN takes the output of CNN as input, and returns sequential predictions of different timesteps. The cascaded CNN-RNN frameworks are often intended for different tasks, rather than image classification. For example, [10, 266, 268] employed CNN-RNN to address the image captioning task, and [273] utilized CNN-RNN to rank the tag list based on the visual importance.

In this chapter, we propose to utilize the cascaded CNN-RNN framework to address a new task, i.e. hierarchical image classification, where we utilize CNN to generate discriminative image features, and utilize RNN to model the sequential relationship of hierarchical labels.

5.2.2 Hierarchical models for image classification

Hierarchical models have been used extensively for image classification. For example, Salakhutdinov et al. [274] presented a hierarchical classification model to share features between categories, and boosted the classification performance for objects with few training examples. Yan et al. [215] presented a hierarchical deep CNN that consists of a coarse component trained over all classes as well as several fine components trained over subsets of classes. Instead of utilizing a fixed architecture for classification, Murdock et al. [275] proposed a regularization method, i.e. Blockout, to automatically learn the hierarchical structure.

Another pipeline to employ hierarchical models tends to improve the classification performance by exploiting the relationship of the categories in the hierarchy. For instance, Deng et al. [264] introduced HEX graphs to capture the hierarchical and exclusive relationship between categories. Ristin et al. [265] utilized Random Forests and proposed a regularized objective function to model the relationship between the categories and subcategories. This type of hierarchical models can not only improve the traditional image classification performance, but also provide an alternative way to utilize the coarse-labeled training data.

In contrast to previous works, our work utilizes RNN to exploit the hierarchical relationship between coarse and fine categories, and aims to adapt the model to address the hierarchical image classification task, in which we simultaneously generate hierarchical labels for the images. Compared with [264, 265] that can only process the pre-computed image features, our proposed CNN-RNN framework can be trained end-to-end.

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

5.3 Hierarchical Image Classification

The goal of our approach is to simultaneously generate hierarchical labels of the images. To this end, we can employ two types of generators: a CNN-based generator and a CNN-RNN generator. Both of them keep the preceding layers of the basic CNN structure except for the last layer.

5.3.1 CNN-based generator

A CNN-based generator aims to generate coarse and fine labels by utilizing the conventional CNN structure. It acts as the general practice to fulfill this specific task. We replace the last layer of conventional CNN with two layers, through which to provide separate supervisory signals for both the coarse categories and fine categories. The two layers can be arranged either in a serial pattern (Figure 5.2: Strategy 1 & 2), or in a parallel pattern (Figure 5.2: Strategy 3).

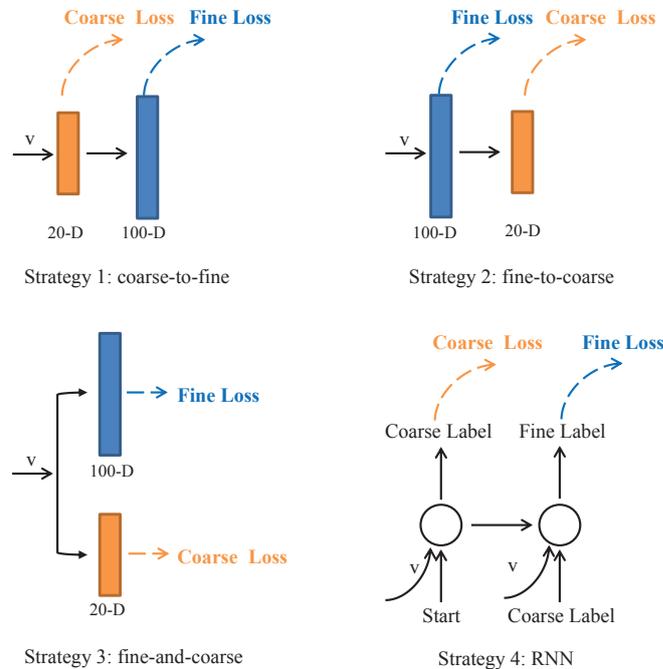


Figure 5.2: The illustration of the four strategies which can jointly train and generate the coarse and fine labels.

5.3 Hierarchical Image Classification

During the training phase, we utilize the softmax loss function to jointly optimize the coarse and fine label predictions, as defined in Eq 5.1.

$$\text{Loss} = -\frac{1}{N} \sum_{i=1}^N \left(\sum_{j=1}^C 1 \{x^i = j\} \log p_j + \sum_{k=1}^F 1 \{y^i = k\} \log p_k \right) \quad (5.1)$$

Where $1 \{\cdot\}$ is the indicator function. N, C, F denote the number of the images, coarse categories, and fine categories, respectively. p_j and p_k are the softmax probabilities of the coarse and fine categories, respectively.

During the inference phase, we can utilize the trained network to determine the coarse and fine labels at the same time.

There are two potential drawbacks for the CNN-based generator: first, it treats the two supervisory signals individually, and does not exploit the relationship between them. Second, when the hierarchy is of variable length, we cannot define a universal CNN-based generator to simultaneously determine the hierarchical labels.

5.3.2 CNN-RNN generator

A CNN-RNN generator determines hierarchical predictions using an architecture where the last layer of CNN is replaced by RNN (Figure 5.2: Strategy 4).

RNN [276] is a class of artificial neural networks where connections between units form a directed cycle, as shown in Figure 5.3. It can effectively model the dynamic temporal behavior of sequences with arbitrary lengths. Long-Short Term Memory (LSTM) [277] is a particular form of a traditional RNN. It extends the cell in a standard RNN by using three gates, including an input gate, a forget gate and an output gate, to accumulate or forget relevant contextual information in its hidden state, as shown in Figure 5.3. These gates enable LSTM to model long-term dependencies in a sequence, and effectively address the gradient vanishing/exploding issues that commonly appear during RNN training [271].

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

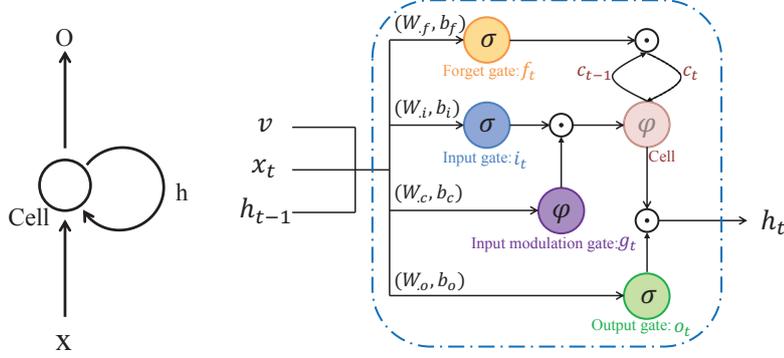


Figure 5.3: The pipeline of RNN(left) and LSTM(right).

In this work, we use LSTM neurons as our recurrent neurons. The definition of the gates and the update of LSTM at the timestep t are as follows:

$$i_t = \sigma(W_{xi}x_t + W_{hi}h_{t-1} + W_{vi}v + b_i) \quad (5.2)$$

$$f_t = \sigma(W_{xf}x_t + W_{hf}h_{t-1} + W_{vf}v + b_f) \quad (5.3)$$

$$o_t = \sigma(W_{xo}x_t + W_{ho}h_{t-1} + W_{vo}v + b_o) \quad (5.4)$$

$$g_t = \varphi(W_{xc}x_t + W_{hc}h_{t-1} + W_{vc}v + b_c) \quad (5.5)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot g_t \quad (5.6)$$

$$h_t = o_t \odot \varphi(c_t) \quad (5.7)$$

Where \odot represents the product operation, σ is the sigmoid function ($\sigma(x) = (1 + \exp(-x))^{-1}$), and φ is the hyperbolic tangent function ($\varphi(x) = \frac{\exp(x) - \exp(-x)}{\exp(x) + \exp(-x)}$). The definition for other symbols are: i_t , f_t , o_t , g_t denote the input gate, forget gate, output gate, and input modulation gate, respectively. x , h , v and c represent the input vector, hidden state, image visual feature, and memory cell, respectively. We propose to impose the image visual feature v at each timestep when updating the LSTM. W and b are the weights and bias that need to be learned.

The goal of our approach is to generate hierarchical labels for images. The labels are ordered in a coarse-to-fine pattern, i.e. coarser labels appear at the front of the list. To this end, we merge the C coarse categories and F fine categories as $C + F$ super categories. For different timesteps, the CNN-RNN generator takes the labels of different levels as input, where the coarser-level labels appear

5.3 Hierarchical Image Classification

at the preceding timesteps. In this way, the coarser-level labels can provide insightful information for the prediction of finer labels. The procedure is shown in Figure 5.4.

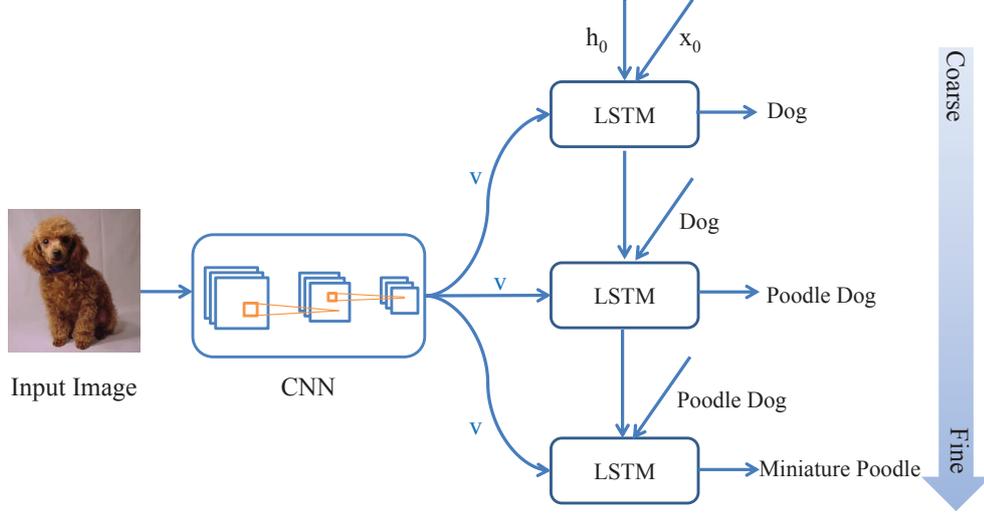


Figure 5.4: The pipeline of the CNN-RNN framework.

During the training phase, the CNN-RNN generator utilizes the groundtruth coarser-level labels as input, and jointly optimizes the coarse and fine predictions, as denoted in Eq 5.8.

$$\text{Loss} = -\frac{1}{N} \sum_{i=1}^N \left(\sum_{t=1}^T \sum_{j=1}^{C+F} 1 \{x_t^i = j\} \log p_j \right) \quad (5.8)$$

During the inference phase, when the groundtruth coarser-level labels are not available, the CNN-RNN generator first predicts the maximum likelihood label for current timestep, i.e. $W_t = \operatorname{argmax}_{W_{t-1}} p(W_{t-1}|I)$, and then utilize the predicted label as the input for the next timestep.

As the CNN-RNN generator defines the super categories, and it equally trains and predicts the super categories, we do not need to design specific networks for the categories of different levels. Therefore, the CNN-RNN generator is robust, and can be employed to generate hierarchical labels of different lengths.

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

5.4 Experiments

We perform our experiments on three well-known datasets: CIFAR-100 [278], ImageNet 2012 [21] and a subset of ImageNet 2010 [265]. These three datasets have provided hierarchical image labels. The characteristics of the three datasets are summarized in Table 5.1.

Table 5.1: The characteristics of the datasets, including the depth of the hierarchy, the number of the coarse categories and fine categories.

Dataset	Depth	Coarse No.	Fine No.
CIFAR-100	2	20	100
ImageNet 2012	1-9	860	1000
Subset of ImageNet 2010	2	143	387

The performance is measured based on the top-1 accuracy. All the experiments are conducted using the Caffe [218] library with a NVIDIA TITAN X card.

The experiments can be divided into two parts. In the first part, we evaluate the performance of hierarchical predictions. In the second part, we investigate the performance of subcategory classification when only a part of the training data is labeled with fine labels while the rest only has coarse labels.

5.4.1 Hierarchical predictions

We evaluate the hierarchical predictions on two widely-used datasets: CIFAR-100 [278] and ImageNet 2012 [21].

5.4.1.1 CIFAR-100

CIFAR-100 contains 100 classes, and each class has 500 training images and 100 test images. These classes are further grouped into 20 superclasses. Therefore, each image comes with two level labels: a fine label (the class to which it belongs) and a coarse label (the superclass to which it belongs). For data preprocessing,

we normalize the data using the channel means and standard deviations. The symbol ‘+’ means a standard data augmentation, i.e. first zero-padded with 4 pixels on each side, and randomly crop 32×32 images from the padded images, or their horizontal reflections.

Evaluation of the hierarchical image classification task.

The CNN-based generator and CNN-RNN generator are considered as two alternative structures to fulfill the hierarchical image classification task. In contrast to CNN-based generator, the CNN-RNN generator can effectively exploit the dependency of the hierarchical labels, and thereby achieving a better classification performance for both the coarse and fine categories. We compare their performance in Table 5.2.

Table 5.2: The comparison of the accuracy for the coarse categories and fine categories. ‘+’ indicates a standard data augmentation (translation/mirroring)

	C100		C100+	
	coarse	fine	coarse	fine
coarse-to-fine	73.88%	58.41%	78.1%	64.16%
fine-to-coarse	75.02%	61.75%	78.16%	65.56%
fine-and-coarse	74.72%	61.8%	77.56%	64.87%
CNN-RNN	80.81%	69.69%	83.21%	72.26%

All of the evaluations in this part are conducted based on the CNN model proposed in [263], because of its high training efficiency and decent performance on CIFAR-100. The CNN structure is shown in Figure 5.5. We employ exactly the same experimental configuration as used in [263].

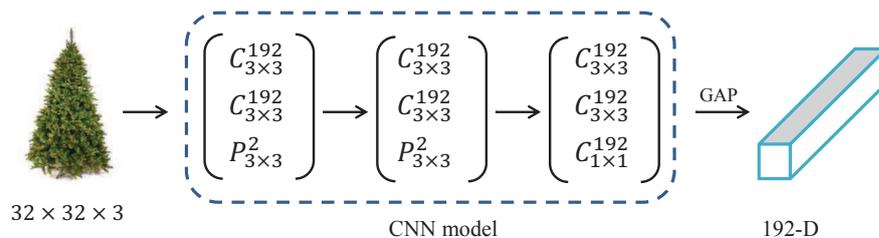


Figure 5.5: The CNN baseline proposed in [263].

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

As can be seen, the CNN-RNN generator can significantly outperform the CNN-based generator, both for the coarse/fine predictions with/without data augmentation. Specifically, for the coarse predictions, the CNN-RNN generator outperforms the CNN-based generator by at least 5.05%, while for the fine predictions, the CNN-RNN generator is even more advantageous, with an improvement of more than 6.7%. This demonstrates that, by exploiting the latent relationship between the coarse and fine categories, RNN can properly address the hierarchical-based task.

Evaluation of the traditional image classification task.

The traditional image classification task consists of classifying images into one pre-defined category, rather than multiple hierarchical categories.

As the CNN-RNN generator can simultaneously generate the coarse and fine labels, in this part, we further compare its performance with ‘coarse-specific’ and ‘fine-specific’ networks. The ‘fine-specific’ network uses the common CNN structure which is specifically employed for the fine category classification. The ‘coarse-specific’ network shares the same preceding layers with the ‘fine-specific’ network, where the last layer is adapted to equal the coarse category number, e.g. 20 for CIFAR-100.

The coarse-specific, fine-specific and CNN-RNN framework can be constructed based on any CNN architecture. To make the comparison more general and convincing, we evaluate the performance on three networks: CNN-7 [263], wrn-28-10 [262] and our proposed wider-Resnet.

For wrn-28-10, we adopt the version with dropout [52], and train the network with larger mini-batch size (i.e. 200), and more iterations (a total of 7×10^4 iterations, and the learning rate dropped at 2×10^4 , 4×10^4 , 6×10^4 iterations). Other experimental configuration follows [262].

The structure of our proposed wider-Resnet is shown in Table 5.3. We adopt the pre-activation residual block as in [279], and train the models for a total of 7×10^4 iterations, with a mini-batch size of 200, a weight decay of 0.0005 and a

momentum of 0.9. The learning rate is initialized with 0.1, and is dropped by 0.1 at 4×10^4 and 6×10^4 iterations.

Table 5.3: The framework of our proposed wider-Resnet.

Group name	Output size	wider-Resnet
conv1	32×32	$3 \times 3, 64$
conv2	32×32	$\begin{bmatrix} 1 \times 1, 256 \\ 3 \times 3, 256 \\ 1 \times 1, 256 \end{bmatrix} \times 3$
conv3	16×16	$\begin{bmatrix} 1 \times 1, 512 \\ 3 \times 3, 512 \\ 1 \times 1, 512 \end{bmatrix} \times 3$
conv4	8×8	$\begin{bmatrix} 1 \times 1, 1024 \\ 3 \times 3, 1024 \\ 1 \times 1, 1024 \end{bmatrix} \times 3$
pool5	1×1	global average pooling

The results on these three datasets are shown in Table 5.4.

Table 5.4: The comparison of the accuracy for the coarse categories and fine categories. ‘+’ indicates standard data augmentation (translation/mirroring)

		C100+	
		coarse	fine
CNN-7 [263]	coarse-specific	82.09%	-
	fine-specific	-	72.03%
	CNN-RNN	83.21%	72.26%
wrn-28-10 [262]	coarse-specific	82.59%	-
	fine-specific	-	74.55%
	CNN-RNN	85.39%	76.23%
wider-Resnet	coarse-specific	85.38%	-
	fine-specific	-	77.97%
	CNN-RNN	88.23%	79.14%

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

We can see that, CNN-RNN can simultaneously generate the coarse and fine labels without developing two separate models, and the accuracy for both categories outperforms the specific networks. Take our proposed wider-Resnet as an example, the CNN-RNN structure increases the coarse and fine accuracy by 2.85% and 1.17% respectively, over the coarse-specific and fine-specific networks. This advantage demonstrates that, by exploiting the latent relationship of the coarse and fine categories, CNN-RNN can help the traditional image classification task.

Our implementation of wrn-28-10 [262] cannot reproduce the original published results, possibly as a result of the differences in the platforms (Torch v.s. Caffe), or the differences in the preprocessing step (pad with reflections of original image v.s. pad with zero). Nevertheless, we can still improve the coarse and fine accuracy by 2.8% and 1.68% respectively, through utilizing the CNN-RNN structure.

Comparison with the state-of-the-art.

We compare our wider-Resnet network, as well as its CNN-RNN variant, with the state-of-the-art, as is shown in Table 5.5.

Through the comparison, we further demonstrate the superiority of the wider networks on CIFAR-100 dataset, as our not-very-deep wider-Resnet network (29 layers) surpasses the performance of the ResNet with super deep layers (1001 layers). In comparison with another wide ResNet [262] with similar depth, wider-Resnet also demonstrates great improvements and remarkably reduces the classification error from 25.45% to 22.03%, under the same platform and pre-processing step.

Overall, our proposed wider-Resnet achieves the best performance over previous works, and wider-Resnet-RNN further increases the state-of-the-art to 20.86%. Nevertheless, we are still seeking to build the CNN-RNN framework on top of future state-of-the-art architectures to boost the classification performance.

Table 5.5: The test error of different methods on the CIFAR-100 dataset with standard data augmentation (translation/mirroring)

Method	C100+
FitNet [280]	35.04%
DSN [281]	34.57%
All-CNN [282]	33.71%
Highway Network [283]	32.39%
APL [284]	30.83%
SReLU [285]	29.91%
BayesNet [286]	27.4%
Fitnet4-LSUV [287]	27.66%
ELU [288]	24.28%
MBA [289]	24.1%
ResNet-110 [227] (according to [290])	27.22%
ResNet-110 (Stochastic Depth) [290]	24.58%
ResNet-164 (Pre-activation) [279]	24.33%
ResNet-1001 (Pre-activation) [279]	22.71%
18-layer + wide RiR [291]	22.90%
FractalNet-20 [292]	23.30%
FractalNet-40 [292]	22.49%
SwapOut V2 [293] (width \times 4)	22.72%
wrn-28-10 [262](our reproduced)	25.45%
wrn-28-10-RNN	23.77%
wider-Resnet	22.03%
wider-Resnet-RNN	20.86%

5.4.1.2 ImageNet 2012

One notable advantage of RNN is that it can generate sequences with variable lengths. To demonstrate this, we investigate the CNN-RNN framework on the widely used ImageNet 2012 dataset [21].

ImageNet is an image dataset organized according to the WordNet hierarchy [294]. It is larger in scale and diversity than other image classification datasets. ImageNet 2012 uses a subset of ImageNet with roughly 1300 images in each of 1000

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

categories. The images are annotated with hierarchical labels of different lengths. In total, there are about 1.2 million training images and 50000 validation images. For all the experiments, we train our model on the training images, and test on the validation images of the ImageNet 2012 dataset.

We utilize the ResNet-152 [227] as our CNN model. For simplicity, pre-trained model weights are kept fixed without fine-tuning. For the RNN model, we use 1000 dimensions for the embedding and the size of the LSTM memory. During the experiments, we first resize all the images to 224×224 pixels and extract the last pooling features utilizing ResNet-152, and then send the features into LSTM for modeling the category dependency.

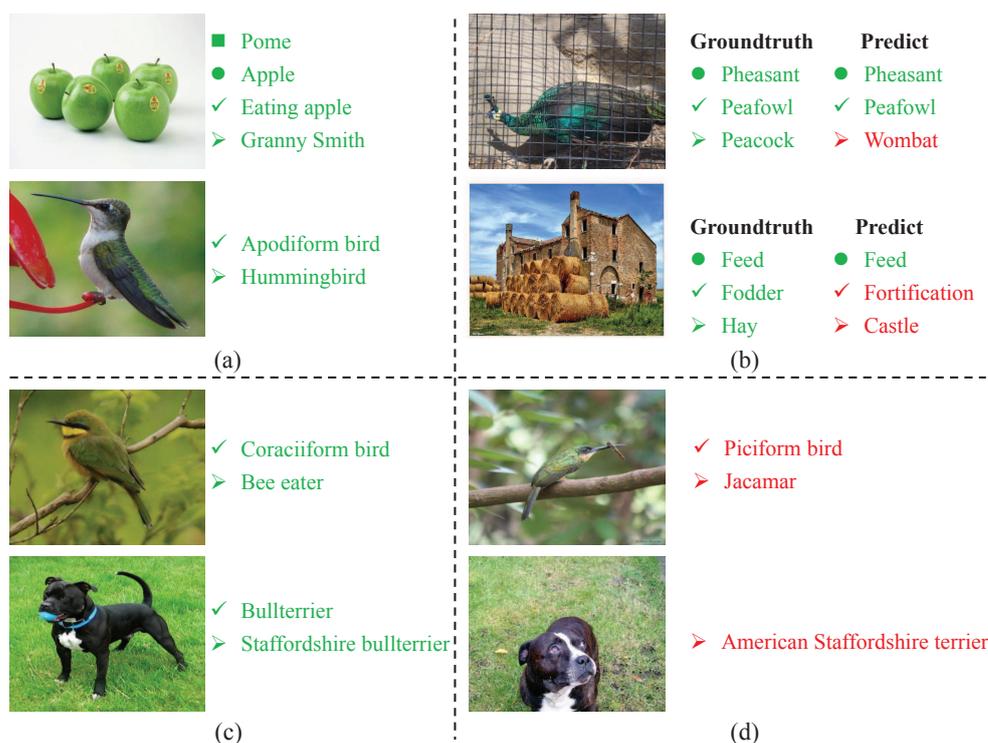


Figure 5.6: The hierarchical predictions of some example images. (a) and (c) show some positive examples. (b) shows the examples with partly wrong predictions, e.g. correct coarse labels & wrong fine labels. (d) shows examples in the same category as (c), but which have totally wrong predictions

Figure 5.6 demonstrates the hierarchical predictions for some example images,

from which we can observe that: First, RNN is able to generate predictions with different lengths, and more specific categories would have more hierarchical labels. Second, the hierarchical labels can describe how the fine categories are evolved from higher level coarse categories, and thus can provide us a better understanding of the objects. Consider for example the upper image in Figure 5.6 (a), we may get confused with the leaf-level label: ‘Granny Smith’. But when the coarse-level labels are provided, we observe that ‘Granny Smith’ is a breed of apple. Third, it may be more difficult to classify images into leaf-level categories than branch-level categories. When we get faulty leaf-level predictions for the given image, we might still learn what the image depicts from the coarse predictions, as shown in Figure 5.6 (b).

5.4.2 From coarse categories to fine categories

In the previous section, we have investigated the hierarchical classification performance of CNN-RNN when all of the coarse and fine labels are available for the training data. However, annotating fine labels for large amounts of training data is quite expensive, especially when it requires expert knowledge. In this subsection, we focus on a scenario in which a part of the training data is annotated with fine labels, while the rest only has coarse labels. This can be viewed as a special case of weakly supervised learning, and has ever been investigated in [265].

We follow the experiment setup of [265], and conduct our experiment on a subset of ImageNet 2010. This dataset particularly selected the classes from ImageNet 2010 that have a unique parent class, and obtained 143 coarse classes and 387 fine ones accordingly. The reduced training set contains 487K images where each coarse class has between 1.4K and 9.8K images, and each fine class has between 668 and 2.4K images. The test set contains 21450 images, and each coarse class has 150 images ¹.

All of the image features are extracted from the VGG-Net [24], as was done for the preliminary experiments in [265].

¹More details about the dataset are available at
http://www.vision.ee.ethz.ch/datasets_extra/mristin/ristinetalcvpr15data.zip

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

Evaluation of the classification performance when all of the training fine labels are available

When all of the coarse and fine labels are available, we can directly train the RNN on the full training set, and evaluate the classification performance on the test set. To better demonstrate the advantage of RNN, we further conduct the training process on a fraction of the training set. In addition, we investigate how much the performance may improve when the coarse labels are provided for the test data, and when we train the CNN-RNN in an end-to-end way, rather than with the off-the-shelf image features. As a comparison with CNN-RNN framework, we also finetune the VGG-Net on the ImageNet 2010 subset. The results are shown in Table 5.6.

Table 5.6: Accuracy for classifying fine labels using the ImageNet 2010 subset described in [265]. RNN: train the RNN with extracted image features from VGG-Net [24]; CNN: finetune the VGG-Net [24] on the ImageNet 2010 subset; CNN-RNN: jointly train the VGG-Net[24] and RNN in an end-to-end pattern; We use the superscript ‘*’ to denote that the coarse labels are provided when predicting the fine labels in the test phase.

	Training Set	Accuracy
NCM [295]	S	66.02%
Multiclass SVM [296]	S	71.67%
RNCMF [265]	S	74.18%
RNN	$0.2S$	75.09%
	$0.4S$	76.17%
	S	77.27%
CNN	S	76.01%
CNN-RNN	S	82%
CNN-RNN*	S	90.69%

We can notice that, training on more data results in a more powerful RNN model, and thus can achieve better performance. Compared with the models trained on parts of the training set, i.e. $0.2S$ and $0.4S$, utilizing the full training set S shows an improvement of 2.18% and 1.1%, respectively. It reveals that, a large training dataset is essential in training the deep models.

In contrast to other methods listed in [265], RNN achieves superior classification performance by inherently exploiting the relationship between the coarse and fine categories. Notably, RNN can deliver better performance even utilizing only 20 percent of the training data.

One additional advantage of the CNN-RNN framework is that it can be trained end-to-end. Compared with the predictions generated with off-the-shelf CNN features, jointly training the CNN and RNN results in a significant improvement, from 77.27% to 82%. It is also much better than directly finetuning the VGG-Net on the ImageNet 2010 subset (82% v.s. 76.01%). When provided the coarse labels for the test images, CNN-RNN achieves an accuracy of 90.69%.

Evaluation of the classification performance when part of the fine labels for training are missing

The training set S in this part are randomly divided into two disjoint sets: S_{coarse} and S_{fine} . S_{coarse} has only the coarse labels, while S_{fine} has both coarse and fine labels. We vary $|S_{\text{fine}}| \in \{0.1|S|, 0.2|S|, 0.5|S|\}$, and for each S_{fine} , we further vary $|S_{\text{coarse}}| \in \{0.1|S|, 0.2|S|, 0.5|S|\}$.

For each training/test configuration, we conduct three evaluations:

- 1) S_{fine} : We train the RNN on S_{fine} , and evaluate on the test set;
- 2) $S_{\text{fine}} + S_{\text{coarse}}^-$: We first train the RNN on S_{fine} , and use it to predict the fine labels of S_{coarse} . In this way, we obtain a new training set S_{coarse}^- , which contains both coarse and (predicted) fine labels. Next, we utilize the S_{fine} and S_{coarse}^- to re-train the RNN, and evaluate on the test set.
- 3) $S_{\text{fine}} + S_{\text{coarse}}^+$: We train the RNN on S_{fine} and S_{coarse}^+ , and evaluate on the test set. S_{coarse}^+ means we utilize the groundtruth fine labels of S_{coarse} .

The results are shown in Figure 5.7.

In general, $S_{\text{fine}} + S_{\text{coarse}}^-$ performs better than S_{fine} , indicating that even some of the fine labels for the training data are missing, the fine category classification can benefit from the CNN-RNN structure.

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK

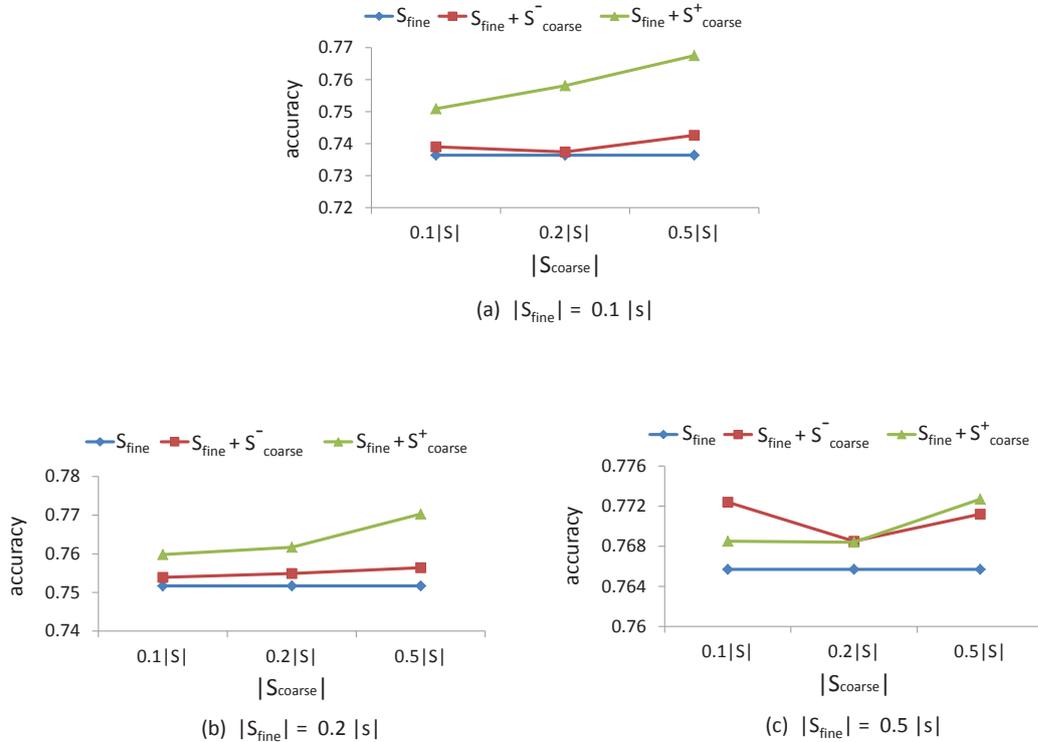


Figure 5.7: The classification performance with different training/test set.

Since the fine labels of S_{coarse} are predicted by the RNN trained on S_{fine} , their accuracy cannot be guaranteed. As a consequence, the second training of RNN may be conducted on a partly wrong labeled dataset. This is particularly severe when $|S_{\text{fine}}|$ is small. As we can see in Figure 5.7(a), when $|S_{\text{fine}}| = 0.1|S|$, the classification hardly benefited from using S_{coarse} when compared to the RNN trained solely on S_{fine} .

On the contrary, when $|S_{\text{fine}}|$ is large, e.g. $|S_{\text{fine}}| = 0.5|S|$, we can achieve a considerable improvement by incorporating S_{coarse} . Notably, when $|S_{\text{fine}}| = 0.5|S|$, $|S_{\text{coarse}}| = 0.1|S|$, $S_{\text{fine}} + S_{\text{coarse}}^-$ even performs slightly better than $S_{\text{fine}} + S_{\text{coarse}}^+$, demonstrating its great potential in weakly supervised classification.

We further compare our method with the NN-H-RNCMF [265], which also attempted to improve the classification by exploiting the hierarchy. We set the amount of coarse-labeled data to $|S_{\text{coarse}}| = 0.5|S|$, and vary the amount of fine-labeled data $|S_{\text{fine}}| \in \{0.1|S|, 0.2|S|, 0.5|S|\}$, and the results are shown in Ta-

ble 5.7. It can be seen that, RNN performs much better than NN-H-RNCMF in all configurations, demonstrating its great potential in exploiting the hierarchical relationship.

Table 5.7: Accuracy in classifying fine categories for the test set. We set the amount of coarse-labeled data to $|S_{\text{coarse}}| = 0.5|S|$.

	$ S_{\text{fine}} $		
	$0.1 S $	$0.2 S $	$0.5 S $
RNCMF [265]	68.49%	70.49%	73.07%
NN-H-RNCMF [265]	69.95%	71.41%	73.43%
RNN	74.26%	75.64%	77.12%

5.5 Conclusion

In this chapter, we proposed to integrate CNN and RNN to accomplish hierarchical classification task. The CNN-RNN framework can be trained end-to-end, and can be built on top of any CNN structures that are primarily intended for leaf-level classification, and further boost the prediction of the fine categories. In addition, we also investigated how the classification would benefit from coarse-labeled training data, which alleviates the professional and expensive manual process of fine-grained annotation.

Currently, it is necessary to have hierarchical labels in the training set, in order to train the RNN. However, this is not available for many small datasets. In the future, we will examine taking advantage of traditional clustering methods towards automatically constructing a hierarchy for the objects, and use CNN-RNN to boost the classification performance for general datasets.

5. CNN-RNN: A LARGE-SCALE HIERARCHICAL IMAGE CLASSIFICATION FRAMEWORK
