

Few-Shot Learning on Cifar-100

Group 14: Hung Cheng, Hsin-Yu Tsai, Chi-Jui Ho, Jing-Cheng Chang

2018.07.04

Abstract

- Goal** 1/5/10 shot learning on Cifar-100
- Dataset** 80 base class: 500/100 for train/test
20 novel class: k-shot images/2000 for train/test
- Approach** Siamese Neural Network [1], Relation Network [2]

Approach

• Siamese Neural Network

• Traditional version

Feature extractor : 5 convolutional layers

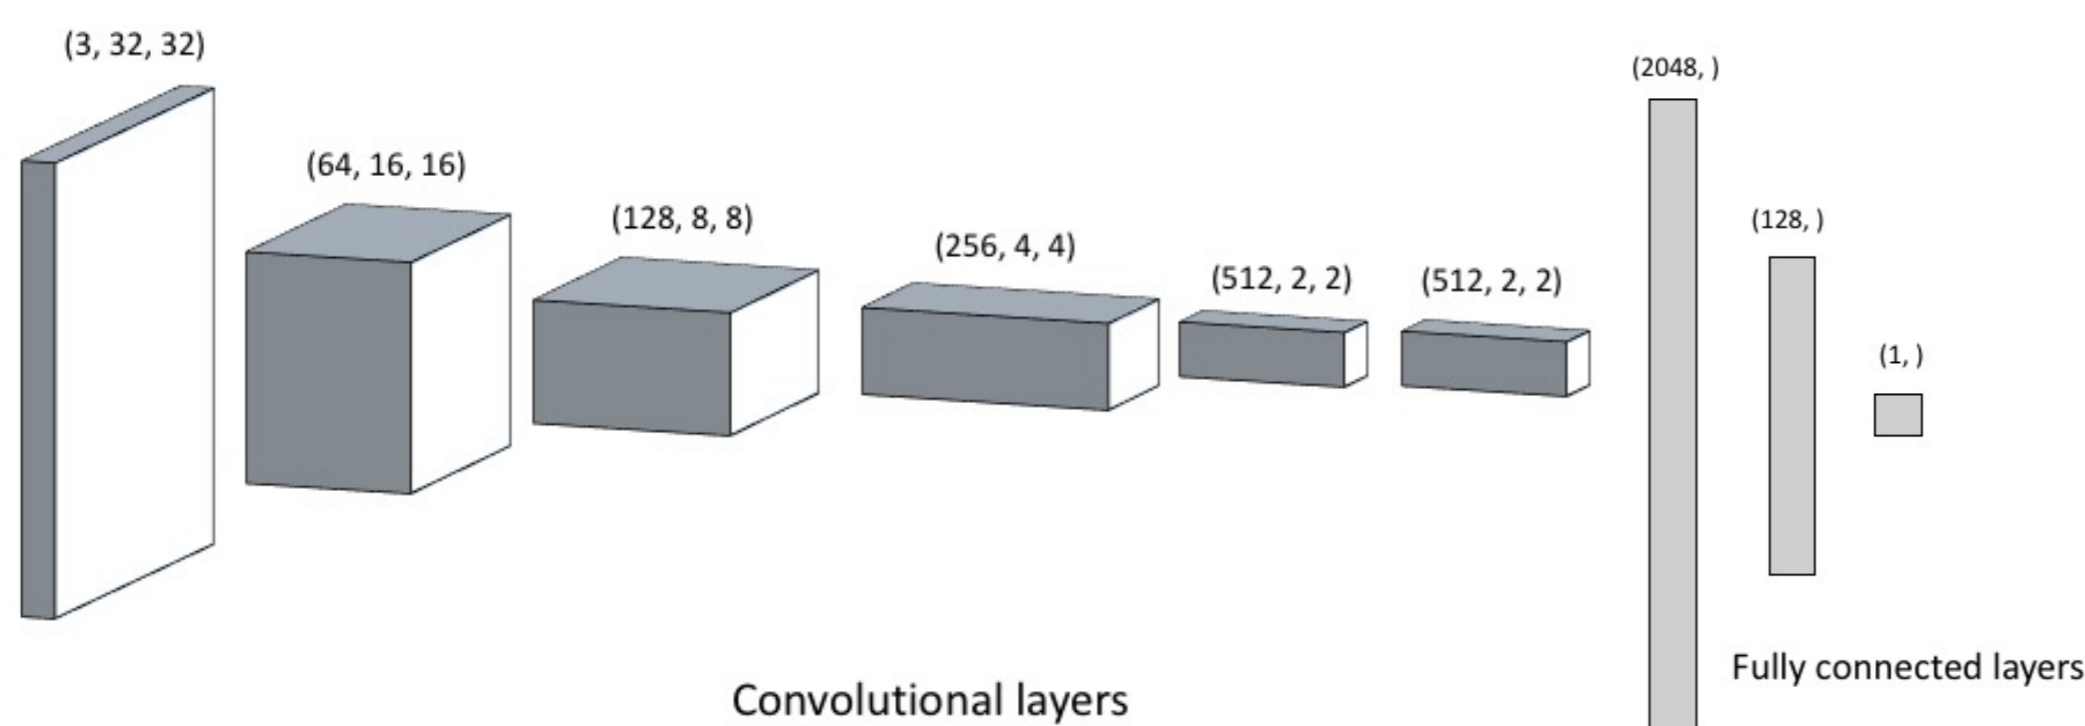
Classifier :

Input : L1 distance of the 2 feature vector

Output : 2 fully connected layer, sigmoid

Loss function : binary cross-entropy

Learn to tell whether 2 images are from same classes



• Our version

Structure is similar to traditional version.

Alternatives : concatenate 2 feature vector, add CNN layer to obtain new relation between 2 feature vector

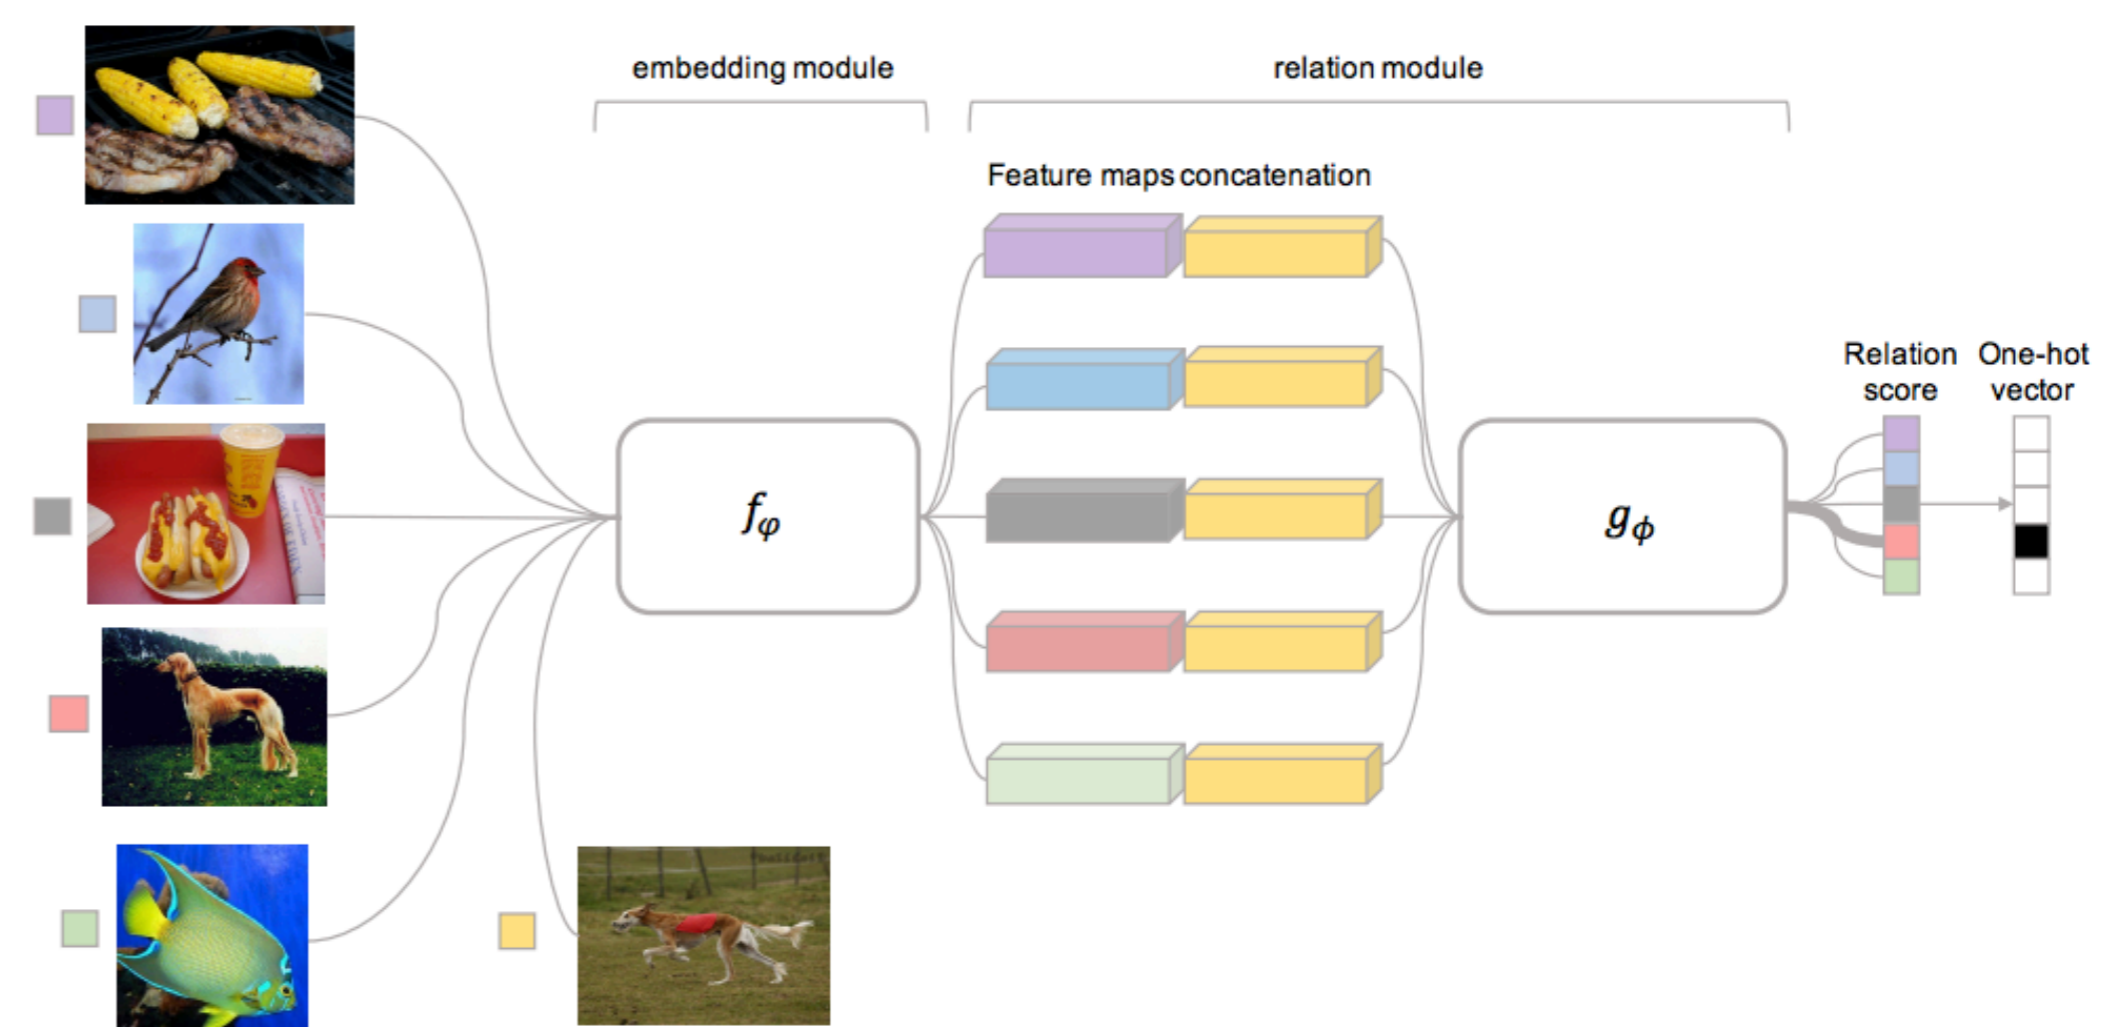
A more effective way to evaluate the correlation.

• Relation Network

Feature embedding : f_ϕ

Relation comparison module : g_ϕ

Loss function : cross-entropy loss



Random sample 20 classes from 80 base classes, each pick k images to stimulate few shot learning.

The model will learn how to compare with 2 feature vector, fix a support set and pick k query image(s) to obtain relations between support set and decide which class has the highest relation score.

Experiment

• Siamese Neural Network

• Comparison accuracy between 1/5/10 shot training

	1-shot	5-shot	10-shot
accuracy	0.3555	0.5125	0.5415

• With/Without Sigmoid at feature extractor output 40 base classes

5-shot	with Sigmoid	without Sigmoid
accuracy	0.2980	0.3420

• Accuracy progress along the augmentation of data pool

classes	20 classes	40 classes	60 classes	80 classes
accuracy	0.2740	0.3795	0.4330	0.4465

• With/Without fine tune

5-shot	without fine tune	with fine tune
accuracy	0.4390	0.5125

• Comparison between traditional method and ours

Accuracy	Traditional method	Our method
validation	0.8070	0.8210
test	0.4205	0.3900

• Relation Network

• Comparison accuracy between 1/5/10 shot training number of query images equals to k

	1-shot	5-shot	10-shot
accuracy	0.2430	0.4455	-

• Accuracy progress compare with validation 50 validation episodes

	1-shot	5-shot	10-shot
validation	0.2760	0.4430	-
test	0.2430	0.4455	-

• With/Without fine tune (with data augmentation)

5-shot	without fine tune	with fine tune
accuracy	0.4000	0.4455

• Data augmentation (without fine tune)

Random horizontal/vertical flip (p=0.1)

Random rotate angle < 15

5-shot	without augmentation	with augmentation
accuracy	0.4215	0.4455

Reference

- [1] Siamese Neural Networks for One-shot Image Recognition, Koch et al., ICML' 15 workshop
- [2] Learning to Compare: Relation Network for Few-Shot Learning, Sung et al., CVPR'18
- [3] Low-shot Visual Recognition by Shrinking and Hallucinating Features, Hariharan et al., ICCV'17