

Neural Network based Classification of Flowers using Transfer Learning

Abhishek
School of Computing Science and
Engineering
Galgotias University
Greater Noida, India
abhishek@galgotiasuniversity.edu.in

Abstract— One of the classical problems in the field of computer vision and machine learning and subsequently deep learning is image classification. While Deep Learning solves the much difficult hurdles like feature extraction and presents us with better optimizations like gradient descent and Adam optimizer, most deep learning models still need a lot of raw computational power to train models on local Graphical Processing Units (GPUs) or Tensor Processing Units (TPUs) in the cloud. All of this computational power is not readily available in all environments and systems and hence the concept of pre-trained models can help to reduce training time by a huge margin. Initial models get trained on large array of GPUs and do feature extraction. The classification part is for the end-user to customize in accordance to the problem at hand and can be completed in very less time.

We tackled the multi-class classification botanical problem of identifying flowers of 5 types, namely, Sunflower, Rose, Dandelion, Daisy, and Tulip. The feature extraction part is done with the model (Google’s Inception-v3) and fully connected softmax layers were trained on local machine on a Nvidia GeForce GTX 950 (with CUDA activated) within 30 minutes time and total steps/epochs were 4000 only. The total number of training images is 3,500 (approx.). The finished model produced results with final test accuracy as 91.9% on new images (N=664).

Keywords—classification, flowers, feature extraction, pre-trained models.

I. INTRODUCTION

Convolutional neural networks (CNNs) opened a door to a new era of classification. Multi-class classification of the order of couple dozen classes is now possible. Using existing CNN model of layering can be efficient as well as robust in classification tasks, but requires a lot of data and computational power to classify images with higher degree of accuracy. This hardware include GPUs, FPGAs, and ASICs like Google’s TPU and IBM TrueNorth [1]. Despite the attractive qualities of CNNs, and despite the relative efficiency of their local architecture, they have still been prohibitively expensive to apply in large scale to high-resolution images. The feature extraction task taken at hand in this paper is done by using Inception-v3 [2].

The network training process is then relatively simpler with a GPU and the final model for a custom use case is ready after training on custom classes for approx. 30 min to 1 hr. This paper does it for five flowers’ 3.500 images with each class having approx. 700 images but it can be extended to any multi-class classification problem with no additional overhead.

II. RELATED WORKS

In [3], Noval general K nearest neighbor classifier GKMNC[4] was used for visual classification. Sparse representation based method[5] for learning and deriving the weights coefficients and FISTA[6] was used for

optimization. CNN-M[7], a pretrained CNN was used for image features extractions then marginal PCA[8] is applied to reduce the dimension of the extracted features. In[9], Alex net[10] model, a deep neural network is used to learn scene image features. During the training phase, series of transformation such as convolution, max pooling, etc are performed to obtained image features. Then two classifier SVM[11] classifier and Softmax[12] classifier are trained using extracted features from the AlexNet model. In[13], Spatial pyramid pooling was used in CNN to eliminate the fixed size input requirements. for this new network structure SPP-net was used, which can generate a fixed length representation regardless of image size. Standard back propagation algorithm was used for training, regardless of the input image. In[14], Kernalized version of Nave bayes Neighbour[15] was used for image classification and SVM[11] classifier was trained on Bag-Of-Features[16] for visual classification. In[17], Extension of the HMAX[18], a four level NN has been used for image classifications. The local filters at first level are integrated into last level complex filters to provide a exible description of object regions. In[19], Nearest neighbor classifier[15] was used for visual classifications. SIFT[20] descriptor to describe shape, HSV[21] values to describe colors and MR filters to describe texture were used.

III. LOGISTICS OF TRANSFER LEARNING

A. Feature Extraction Task

One of the most computationally intensive tasks in neural network approach is feature extraction. The transfer learning model aims to make this task independent of the use case. One of the ways it does that is to exhaustively train (primary) a custom CNN architecture on millions of images and labels.

A number of such models are already available e.g. AlexNet[22], VGGNet[23], Google LeNet/Inception-v1[24], ResNet[25].

B. Classification Task

The model is then exported and then subjectively trained (secondary) for a given use case. This training will not require an elaborate dataset, nor will it require powerful computational hardware. The model does not perform feature extraction again during this secondary training but trains and stores weights for new use case among the softmax layers.




Fig. 1. Transfer Learning Pipeline.

IV. INCEPTION-V3 LAYERS [2]

Convolution layer - The model is looking for features in the picture, however, it does not know where they are, so it tries the filter everywhere. The filter is simply a mathematical vector, where it multiplies the value in the image (from 0 to 1) with the value in the convolution.




Fig. 2. Convolution process schematic.

Pooling layer — AvgPool and MaxPool – Reduce the sample size while keeping mathematical integrity of the data.




Fig. 3. Pooling sublayers schematic.

AvgPool sublayer - Average pooling takes the values to process, and takes the average of the values. It adds up all the values, and divides them by the number of values.

MaxPool sublayer - Max pooling takes the values and replaces it with the largest value.

Concatenate layer - combine the results at the end to create a new sample. This step involves dropping all the negative values. We want to filter out the irrelevant parts of the image, so whenever a pixel's value is less than 0, it's swapped out for a 0. Inception-v3 uses ReLu (Rectified Linear Units).




Fig. 4. Concatenate process schematic.

V. DATASET

Inception-v3 is trained for the ImageNet [26] Large Visual Recognition Challenge using data from 2012. ImageNet has

over 14 million images from over 1000 classes and is extensively used for computer vision research.

For this paper, 3,670 images are used to train the fully-connected and softmax layers over 4000 epochs.

All images need to be of clear resolution greater than 800x600 for better results. Images also need to be at least 16-bit coloured. The images can be scraped from google images using chrome extensions or python scripts written for the same.

Each image is resized to 299 x 299 x 3 dimensions because to train a CNN using transfer learning, image input size to CNN must be same as the input size given to the original model [2].

TABLE I. DATASET CLASS LABEL DISTRIBUTION

Class Label	N = No. of images in dataset
Daisy	633
Dandelion	898
Roses	641
Sunflowers	699
Tulips	799
Total	3,670

^a. All images sourced from open-license sources on Google Images.

VI. TRAINING METHODOLOGY

The learning method used in this experiment is supervised learning [27]. The network is trained with stochastic gradient utilizing the TensorFlow [28] distributed machine learning system using a Nvidia GeForce 950 GTX GPU with batch size 100 for 4,000 epochs. Training set has 90 percent images and Evaluation set has two subsets, test subset has 10 percent of images and validation subset has 10 percent of images, both amounting to approx. 367 images each. Training learning rate = 0.01, training interval = 10. These hyperparameters are customizable in *retrain.py* file. Training accuracy was 89.7% at the beginning of the training process and starts to increase, after completion of all training steps it reached to 98.0%. Validation accuracy was 75.8% during initiation of training and validation process and final validation accuracy was 93.0%. Final training accuracy was 91.9.

VII. RESULTS AND COMPARISONS

The model is able to classify class label images with high precision.

Below image of dandelion tested produces the following results –

Evaluation time (1-image): 1.387s
dandelion (score=0.99877)
daisy (score=0.00097)
tulips (score=0.00018)
sunflowers (score=0.00007)
roses (score=0.00001)



Fig. 5. Sample test image (Dandelion class).

Another image test results – Evaluation time (1-image): 1.231s, sunflowers (score=0.93991), dandelion (score=0.03731), daisy (score=0.01330), tulips (score=0.00671), roses (score=0.00277).



Fig. 6. Sample test image (Sunflower class)

CONCLUSION

In this paper, the classification layers (fully-connected and softmax) of pre-trained Inception-v3 model was re-trained successfully by implementing transfer learning technique. The model yields a final test accuracy of 91.9 percent on 5 classes of flower images dataset. Further this model can be used to retrain for much more classes within small time and with reasonable computational hardware to classify plants and other flowers' images as well as other objects.

REFERENCES

- [1] Y. LeCun, "1.1 Deep Learning Hardware: Past, Present, and Future," 2019 IEEE International Solid-State Circuits Conference - (ISSCC), San Francisco, CA, USA, 2019, pp. 12-19.
- [2] Christian Szegedy and Vincent Vanhoucke and Sergey Ioffe and Jonathon Shlens and Zbigniew Wojna, "Rethinking the Inception Architecture for Computer Vision" in 1512.00567- arXiv, 2015
- [3] Q. Liu, A. Puthenputhussery, and C. Liu, Novel general knn classifier and general nearest mean classifier for visual classification," in Image Processing (ICIP), 2015 IEEE International Conference on. IEEE, 2015, pp. 1810-1814
- [4] J. M. Keller, M. R. Gray, and J. A. Givens, A fuzzy k-nearest neighbor algorithm," IEEE transactions on systems, man, and cybernetics, no. 4, pp. 580-585, 1985.
- [5] J. A. Tropp, Greed is good: Algorithmic results for sparse approximation," IEEE Transactions on Information theory, vol. 50, no. 10, pp. 2231-2242, 2004.
- [6] A. Beck and M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems," SIAM journal on imaging sciences, vol. 2, no. 1, pp. 183-202, 2009.

- [7] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, Return of the devil in the details: Delving deep into convolutional nets," arXiv preprint arXiv:1405.3531, 2014.
- [8] I. Jolliffe, Principal component analysis," in International encyclopedia of statistical science. Springer, 2011, pp. 1094-1096.
- [9] J. Sun, X. Cai, F. Sun, and J. Zhang, Scene image classification method based on alex-net model," in Informative and Cybernetics for Computational Social Systems (ICSS), 2016 3rd International Conference on. IEEE, 2016, pp. 363-367.
- [10] A. Krizhevsky, I. Sutskever, and G. E. Hinton, Imagenet classification with deep convolutional neural networks," in Advances in neural information processing systems, 2012, pp. 1097-1105.
- [11] C. Cortes and V. Vapnik, Support-vector networks," Machine learning, vol. 20, no. 3, pp. 273-297, 1995.
- [12] N. M. Nasrabadi, Pattern recognition and machine learning," Journal of electronic imaging, vol. 16, no. 4, p. 049901, 2007.
- [13] K. He, X. Zhang, S. Ren, and J. Sun, Spatial pyramid pooling in deep convolutional networks for visual recognition," in European conference on computer vision. Springer, 2014, pp. 346-361.
- [14] T. Tuytelaars, M. Fritz, K. Saenko, and T. Darrell, The nbnn kernel," in Computer Vision (ICCV), 2011 IEEE International Conference on. IEEE, 2011, pp. 1824-1831.
- [15] K. P. Murphy et al., Naive bayes classifiers," University of British Columbia, vol. 18, 2006.
- [16] Z. S. Harris, Distributional structure," Word, vol. 10, no. 2-3, pp. 146-162, 1954.
- [17] C. Theriault, N. Thome, and M. Cord, Extended coding and pooling in the hmax model," IEEE Transactions on Image Processing, vol. 22, no. 2, pp. 764-777, 2013.
- [18] M. Riesenhuber and T. Poggio, Hierarchical models of object recognition in cortex," Nature neuroscience, vol. 2, no. 11, p. 1019, 1999.
- [19] M.-E. Nilsback and A. Zisserman, A visual vocabulary for flower classification," in Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on, vol. 2. IEEE, 2006, pp. 1447-1454.
- [20] D. G. Lowe, Object recognition from local scale-invariant features," in Computer vision, 1999. The proceedings of the seventh IEEE international conference on, vol. 2. Ieee, 1999, pp. 1150-1157.
- [21] A. R. Smith, Color gamut transform pairs," ACM Siggraph Computer Graphics, vol. 12, no. 3, pp. 12-19, 1978.
- [22] Krizhevsky, Alex & Sutskever, Ilya & Hinton, Geoffrey. (2012). ImageNet Classification with Deep Convolutional Neural Networks. Neural Information Processing Systems. 25. 10.1145/3065386.
- [23] Simonyan, Karen & Zisserman, Andrew. (2014). Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv 1409.1556.
- [24] C. Szegedy, Wei Liu, Yangqing Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Jun. 2015.
- [25] K. He, X. Zhang, S. Ren and J. Sun, "Deep Residual Learning for Image Recognition," 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, 2016, pp. 770-778.
- [26] J. Deng, W. Dong, R. Socher, L. Li, Kai Li and Li Fei-Fei, "ImageNet: A large-scale hierarchical image database," 2009 IEEE Conference on Computer Vision and Pattern Recognition, Miami, FL, 2009, pp. 248-255.
- [27] S. B. Kotsiantis, I. Zaharakis, and P. Pintelas, Supervised machine learning: A review of classification techniques," Emerging artificial intelligence applications in computer engineering, vol. 160, pp. 3-24, 2007.
- [28] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mane, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viegas, O. Vinyals, P. Ward, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org