

Learning Abstract Classes using Deep Learning

Sebastian Stabinger
University of Innsbruck,
Institute of Computer Science
Technikerstrasse 21a
Innsbruck, Austria
Sebastian.
Stabinger@uibk.ac.at

Antonio
Rodríguez-Sánchez
University of Innsbruck,
Institute of Computer Science
Technikerstrasse 21a
Innsbruck, Austria
Antonio.Rodriguez-
Sanchez@uibk.ac.at

Justus Piater
University of Innsbruck,
Institute of Computer Science
Technikerstrasse 21a
Innsbruck, Austria
Justus.Piater@uibk.ac.at

ABSTRACT

Humans are generally good at learning abstract concepts about objects and scenes (e.g. spatial orientation, relative sizes, etc.). Over the last years convolutional neural networks have achieved almost human performance in recognizing concrete classes (i.e. specific object categories). This paper tests the performance of a current CNN (GoogLeNet) on the task of differentiating between abstract classes which are trivially differentiable for humans. We trained and tested the CNN on the two abstract classes of horizontal and vertical orientation and determined how well the network is able to transfer the learned classes to other, previously unseen objects.

Categories and Subject Descriptors

I.2.10 [Vision and Scene Understanding]: Shape; I.5.4 [Applications]: Computer vision; I.4.8 [Scene Analysis]: Shape

General Terms

Experimentation, Performance

Keywords

Deep Learning, Convolutional Neural Networks, Visual Cortex, Abstract Reasoning

1. INTRODUCTION

Deep learning methods have gained interest from the machine learning and computer vision research community over the last years because these methods provide exceptional performance in classification tasks. Especially Convolutional Neural Networks (CNNs) — first introduced in 1989 by LeCun et al. [9] — have become popular for object classification. CNNs were more widely used after the deep CNN from Krizhevsky et al. [8] outperformed the state of the art methods in ILSVRC12 [10] by a wide margin in 2012.

Convolutional neural networks consist of multiple layers of nodes, also called neurons. One important layer type is the convolutional layer from which the networks obtain their name. In a convolutional layer the responses of the nodes depend on the convolution of a region of the input image with a kernel. Additional layers introduce non-linearities, rectification, pooling, etc. The goal of training a CNN lies in optimizing the network weights using image-label pairs to best reconstruct the correct label given an image. During testing the network is confronted with novel images and expected to generate the correct label. The network is trained by gradient descent. The gradient is calculated by back-propagation of labeling errors. The general idea of CNNs is to automatically learn the features needed to distinguish classes and generate increasingly abstract features as the information moves up the layers.

Since CNNs are very popular at the moment and perceived — in parts of the computer vision community — as obtaining human like performance we wanted to test their applicability on visual tasks slightly outside the mainstream which are still trivially solvable by humans. We chose to learn simple abstract classes using a standard CNN not because we assume that they will perform better on the tasks than other, possibly much simpler methods, but because we want to gain insights into CNNs and how they perform on tasks which can be solved trivially by humans. We will mainly try to give insight into the amount of training images needed and how well the classifier generalizes to previously unseen shapes representing the same abstract concepts. Until now, most of the classes used for training and testing convolutional neural networks were concrete (e.g. detecting classes of objects, animal species in an image, ...). One notable exception is the work by Gülçehre et al.[6] who trained a CNN to recognize whether multiple presented shapes are the same. This is in essence a training on two abstract classes.

The problems presented by Bongard [1] inspired us to do the research presented in this paper. Foundalis [4] gives a good introduction to these problems and presents a system intended to solve them computationally. He used other methods than deep learning though. Previously, Fleuret et al.[3] compared human performance to classical machine learning methods (Adaboost on decision stumps and Support Vector Machines) on classification tasks. The classes used were similar in spirit to Bongard problems (e.g. object similarity, relative position, ...). Problems of similar nature are also

often given at aptitude and intelligence tests to measure the ability for abstract, non-verbal reasoning.

We decided to use the two classes *horizontal* and *vertical* for our classification experiments since they are abstract, visually unambiguous, easy to differentiate by humans, and easily transfer to very different shapes. The goal is to learn a classifier that can distinguish between horizontally and vertically oriented structures and can transfer this knowledge to previously unseen objects and shapes. These are the first experiments exploring the representational capabilities of current deep learning systems regarding abstract classes.

2. MATERIALS AND METHODS

For this paper we used the convolutional neural network GoogLeNet as presented by Szegedy et al. [11]. It won in a number of categories in ILSVRC14. We slightly adapted the implementation provided with the Caffe [7] deep learning framework to our task (i.e. differentiating horizontal from vertical shapes).

For all experiments we started with an initial learning rate of 0.01 and use ADAGRAD [2] to adapt the learning rate over time. We trained the CNN for 1000 iterations. At this point the loss was so small (< 0.01) that no further meaningful improvement was possible. The CNN was trained 10 times on sets of different randomly generated images to judge the mean accuracy as well as the variance for different numbers of training images. All the graphs in this paper show the mean accuracy as blue dots and 90% of all measurements fall within the shaded area (see Figure 2 for an example). The test set contained 250 images per class. The reported accuracy is the proportion of correctly classified images of this test set.

3. EXPERIMENTS

To test how well GoogLeNet can generalize abstract classes to different shapes or different renderings we use the following procedure: We train GoogLeNet without pre-training on a dataset consisting of the two classes “horizontal” and “vertical”. We then test the performance of the net on a test set containing the same two classes but represented by different shapes or rendered differently (e.g. outline of the shape versus a filled representation).

We are interested whether the CNN can distinguish the two classes and the amount of training images we need to obtain satisfactory results.

3.1 Learning on Rectangles, Testing on Ellipses

We randomly generated vertically and horizontally oriented, filled rectangles on a white background for training the network. We tested on randomly generated, vertically or horizontally oriented, filled ellipses (Figure 1).

Figure 2 shows the accuracy of the net after 1000 iterations in relation to the number of training images used per class. The CNN was able to learn and generalize the two classes more or less perfectly with about 100 training images per class. To our surprise, even 10 images per class result in a mean accuracy of about 90%.




Figure 1: Examples of randomly generated, filled rectangles and ellipses (horizontal and vertical class)

As can be expected, the variance is higher for fewer training images. One has to assume that some images are better representations of the classes than others. Since the images are randomly generated the quality of the whole dataset will vary more for fewer images.




Figure 2: Learned on filled rectangles, tested on filled ellipses. Achieved accuracy depending on the number of training images per class.

3.2 Learning on Outline, Testing on Filled

To test how sensitive the network is regarding different representations of the same shape we trained the network for the “horizontal” and “vertical” classes on outlines of rectangles (Figure 3). We used the filled rectangles from Figure 1 for testing.

Figure 4 shows that the network has much bigger problems to generalize from outlines to filled versions of the same shape than it has at generalizing from one filled shape to another (i.e. from rectangle to ellipse). In addition, the vari-




Figure 3: Examples of randomly generated rectangle outlines (horizontal and vertical class)

ance does not decrease with the number of training images. This might indicate that the network is learning features that do not capture the abstract concept but specific information for outlines instead (i.e. overfitting to the training set).




Figure 4: Learned on rectangle outlines, tested on filled rectangles. Achieved accuracy depending on the number of training images per class.




Figure 5: Learned on rectangle and ellipse outlines, tested on filled rectangles. Achieved accuracy depending on the number of training images per class.

If a CNN is able to learn abstract concepts one can reason that adding another shape outline to the training set will improve the accuracy. By this we force the net to learn a more abstract concept. To test this hypothesis we added horizontally and vertically oriented ellipse outlines to the training images and again tested the accuracy on the filled rectangles. As can be seen in Figure 5, the addition of ellipse outlines improved the performance on the filled rectangle test set slightly.

3.3 Random Shapes




Figure 6: Examples of randomly generated shapes (horizontal and vertical class)

We performed the last set of experiments on random shapes (Figure 6) which we created with an adapted version of the SVRT framework presented by Fleuret et al. [3].

The network is able to learn the two abstract classes with these highly varied shapes (Figure 7) and is also able to transfer the knowledge from outlines to filled shapes (Figure 9).




Figure 7: Learned on random outlines, tested on random outlines. Achieved accuracy depending on the number of training images per class.

Figure 10 shows that the network trained on random shape outlines even performs better at detecting the orientation of filled rectangles than the network trained on similar data sets (Figure 4 and Figure 5)

As a final experiment we looked at how well the random shapes generalize to other, textured random shapes. We used the likely most difficult test set where the texture orientation was orthogonal to the orientation of the shape. Figure 11 shows examples of this class. The results (Figure 12) indicate that more training examples lead to extreme variance in the results. The performance of the network varies from perfect accuracy to pure guessing. In addition, nothing during the training phase indicates how well the network will perform on the test set (e.g. lower loss during training is uncorrelated to the performance on the test set).

4. CONCLUSION

We showed that a state of the art convolutional neural network is able to learn abstract classes and transfer that information to other, previously unseen, shapes. But it is also apparent that the current networks are sensitive to the used training and testing data regarding how well the transfer of knowledge works.

Probably the best example is the training on filled rectan-



Figure 8: Examples of randomly generated, filled shapes (horizontal and vertical class)




Figure 9: Learned on random outlines, tested on random, filled shapes. Achieved accuracy depending on the number of training images per class.




Figure 12: Learned on random outlines, tested on texture filled random shapes. Achieved accuracy depending on the number of training images per class.




Figure 10: Learned on random outlines, tested on filled rectangles. Achieved accuracy depending on the number of training images per class.

gles and testing on filled ellipses in comparison to training on rectangle outlines and testing on filled rectangles. It was unclear before doing experiments that GoogLeNet will perform much better on the first task than the second one. Humans are in general much less affected by such representational differences and perform well on highly variable data sets as can be seen in problem sets used for measuring non-verbal abstract reasoning or the Bongard problems.

Of course humans as well as animals already have pre-training before encountering such tasks. Pre-training in the form of previously learned concepts as well as the optimization that occurred during evolution and manifests itself in the organization of the brain. A CNN is missing this information. We think therefore that pre-training of networks on the right data set will be paramount to increase the performance on more abstract tasks. The results of Gülçehre et al.[6] also




Figure 11: Examples of randomly generated, textured shapes (horizontal and vertical class)

point into this direction.

5. REFERENCES

- [1] M. M. Bongard. *Pattern Recognition*. Spartan Books, 1970.
- [2] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.
- [3] F. Fleuret, T. Li, C. Dubout, E. K. Wampller, S. Yantis, and D. Geman. Comparing machines and humans on a visual categorization test. *Proceedings of the National Academy of Sciences*, 108(43):17621–17625, 2011.
- [4] H. E. Foundalis. Phaeaco: a cognitive architecture inspired by bongard’s problems. *Unpublished Ph. D. Thesis, Department of Computer Science and the Cognitive Science Program. Indiana University, Bloomington, IN*, 2006.
- [5] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *International conference on artificial intelligence and statistics*, pages 249–256, 2010.
- [6] Ç. Gülçehre and Y. Bengio. Knowledge matters: Importance of prior information for optimization. *arXiv preprint arXiv:1301.4083*, 2013.
- [7] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [8] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [9] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- [10] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, pages 1–42, April 2015.

- [11] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *arXiv preprint arXiv:1409.4842*, 2014.

6. APPENDIX — GOOGLNET

In this appendix we will give a brief introduction to GoogLeNet, the convolutional neural network used for the experiments in this paper.

6.1 Inception Module

An inception module is a small two layer network which is repeated many times. It is used as the main building block of GoogLeNet. Figure 13 shows a graphical representation of an inception module.

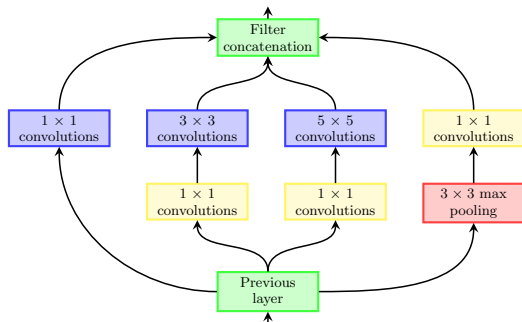


Figure 13: Overview of an inception module. Adapted from Szegedy et al. [11]. The 1×1 convolutions for dimensionality reduction are shown in yellow.

An inception module computes convolutions with different receptive field sizes (1×1 , 3×3 , 5×5) and 3×3 max pooling in parallel and concatenates all the responses to produce the output to the next layer. Since this would lead to excessive amounts of parameters, 1×1 convolutions are used for dimensionality reduction.

6.2 GoogLeNet

GoogLeNet consist of a stack of nine inception modules. There are three points at which softmax is being used to calculate the loss of the network. One at the end of all nine inception modules and two after the third and sixth inception module. The reasoning behind using middle layers to calculate an error function is to promote better discrimination in lower layers and to calculate a better gradient signal. Both is needed since we are dealing with a very deep network with 27 layers. We refer to Figure 3 in the paper by Szegedy et al. [11] for a more detailed description of the layer structure of GoogLeNet.

6.3 Specifics of the Implementation

The implementation in the Caffe framework differs slightly from the network presented by Szegedy et al. [11].

Stochastic gradient descent with momentum is used to update the weights and the Xavier algorithm as presented by Glorot et al. [5] is used for initializing the weights.