



Transfer Learning on Object Classification

Introduction

The power of image analysis has been extended by powerful deep learning algorithms. While very flexible, these algorithms require large amounts of data to learn a task, such as object classification. Transfer Learning (TL) offers an opportunity to increase performance with less data.




Fig 1: shows the idea behind Transfer Learning, that beforehand transferring knowledge from another domain improves performance in multiple ways

Hypothesis

The underlying idea behind TL is that the shallow layers in the neural network corresponds to general, and thus transferable, features. By freezing these layers whilst retraining the later layers one may achieve a higher learning rate.

Project

To do this we are conducting a quantitative study of transfer learning, specifically in regards to how many layers one should freeze. If there are general and specific layers, one may be able to determine an optimal amount of layers to freeze for a dataset.

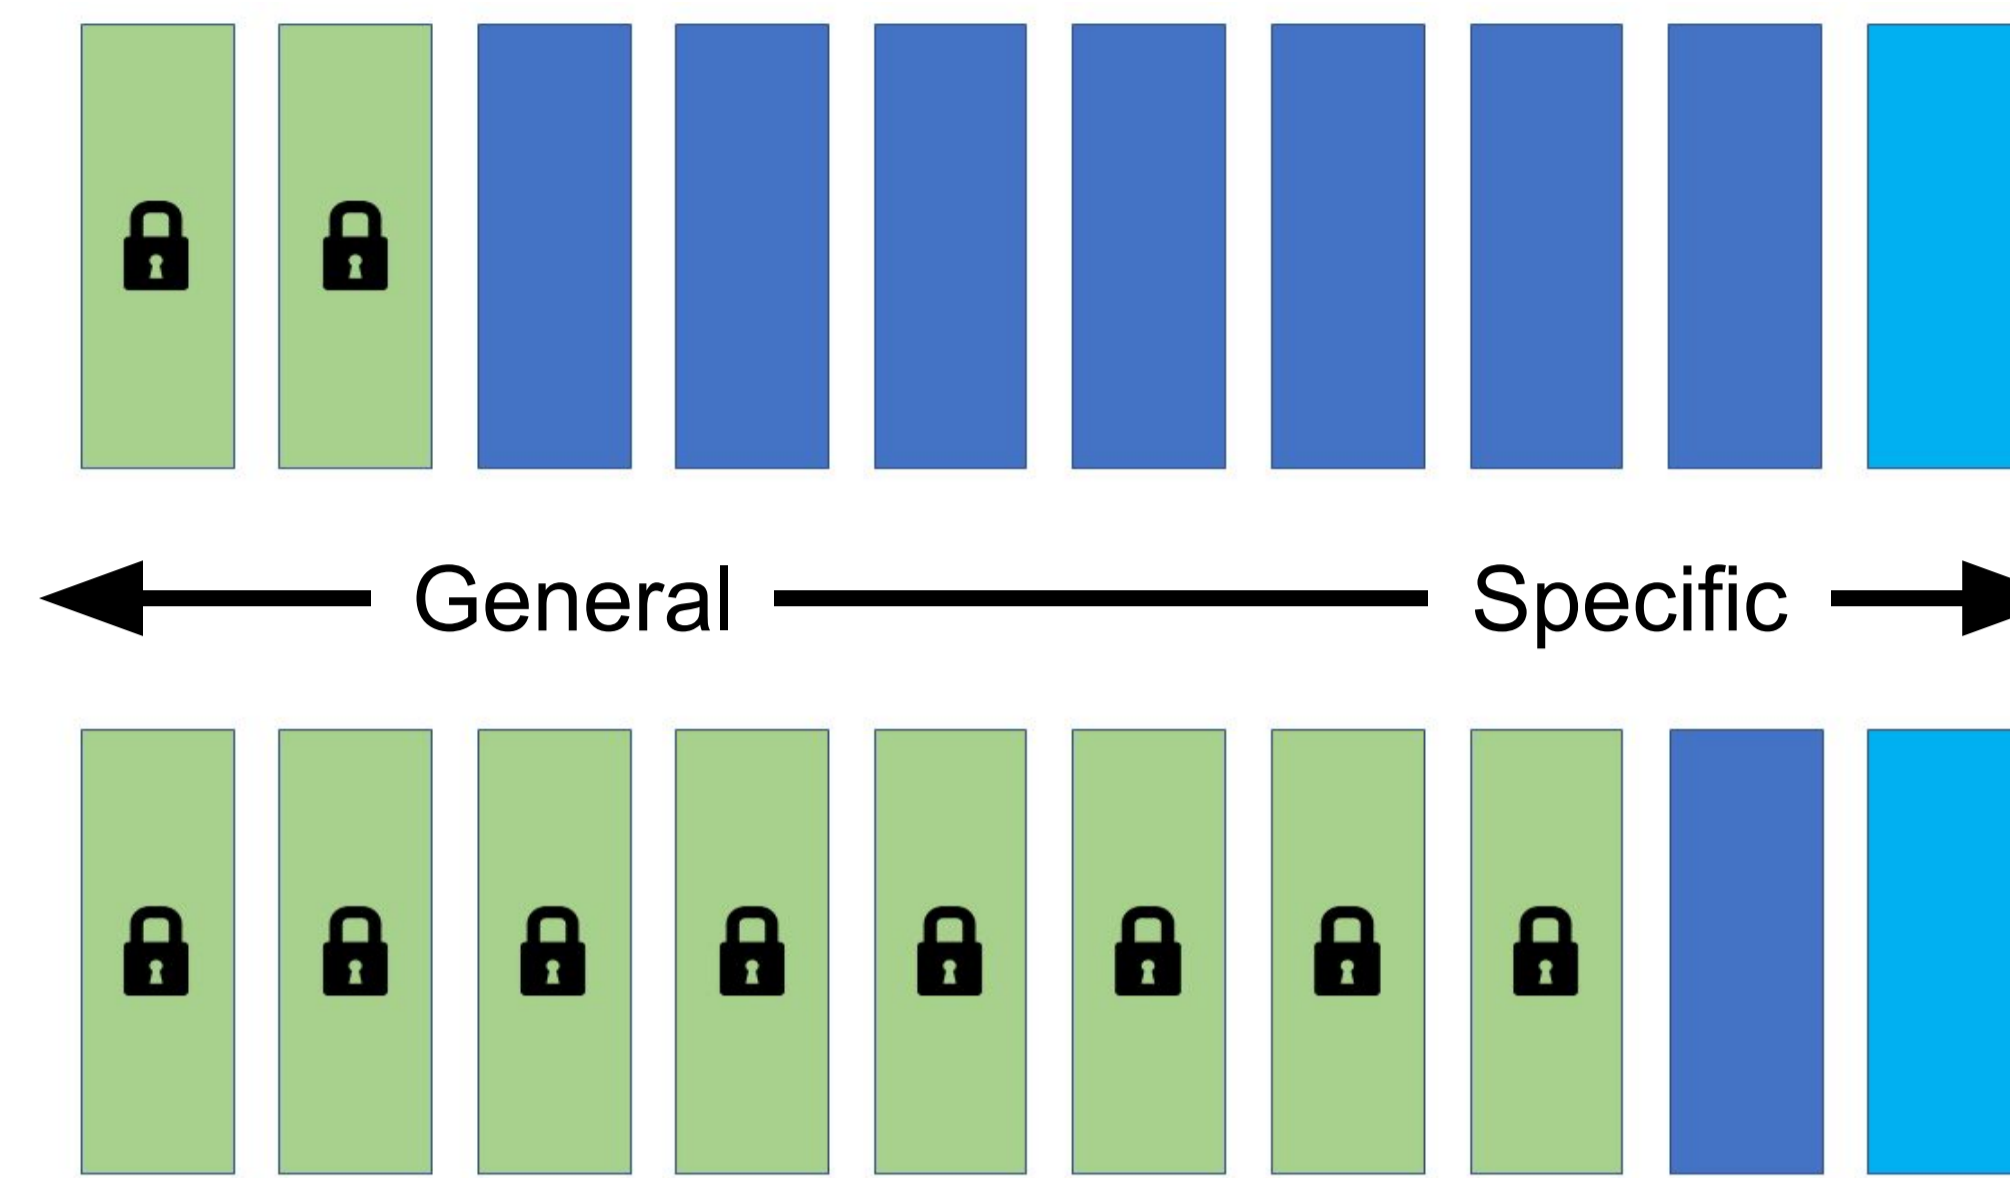


Fig 2: shows the key concept of our method and the underlying theoretical hypothesis. Green layers are frozen, blue retrainable, and cyan both open and untrained.

Method

The transfer is done by taking a neural network trained on a different domain, the **source domain**, and retraining this network on a new domain, the **target domain**. The network layers can be further trained, or kept constant after transfer. Varying the number of frozen layers and dataset size may highlight when TL works well.

Result

In figure 3 we can see the results of applying TL with different amounts of frozen layers and data, both as the actual accuracy and the improvements from TL. The result suggests that there's a correlation between the amount of data and the number of frozen layers which determines the efficiency of the transfer.

Further research

- More detailed studies could answer how well different domains transfer to each other.
- In order to be able to estimate how many layers one should freeze beforehand based on amount of data it is necessary to be able to determine the proximity of the domains.

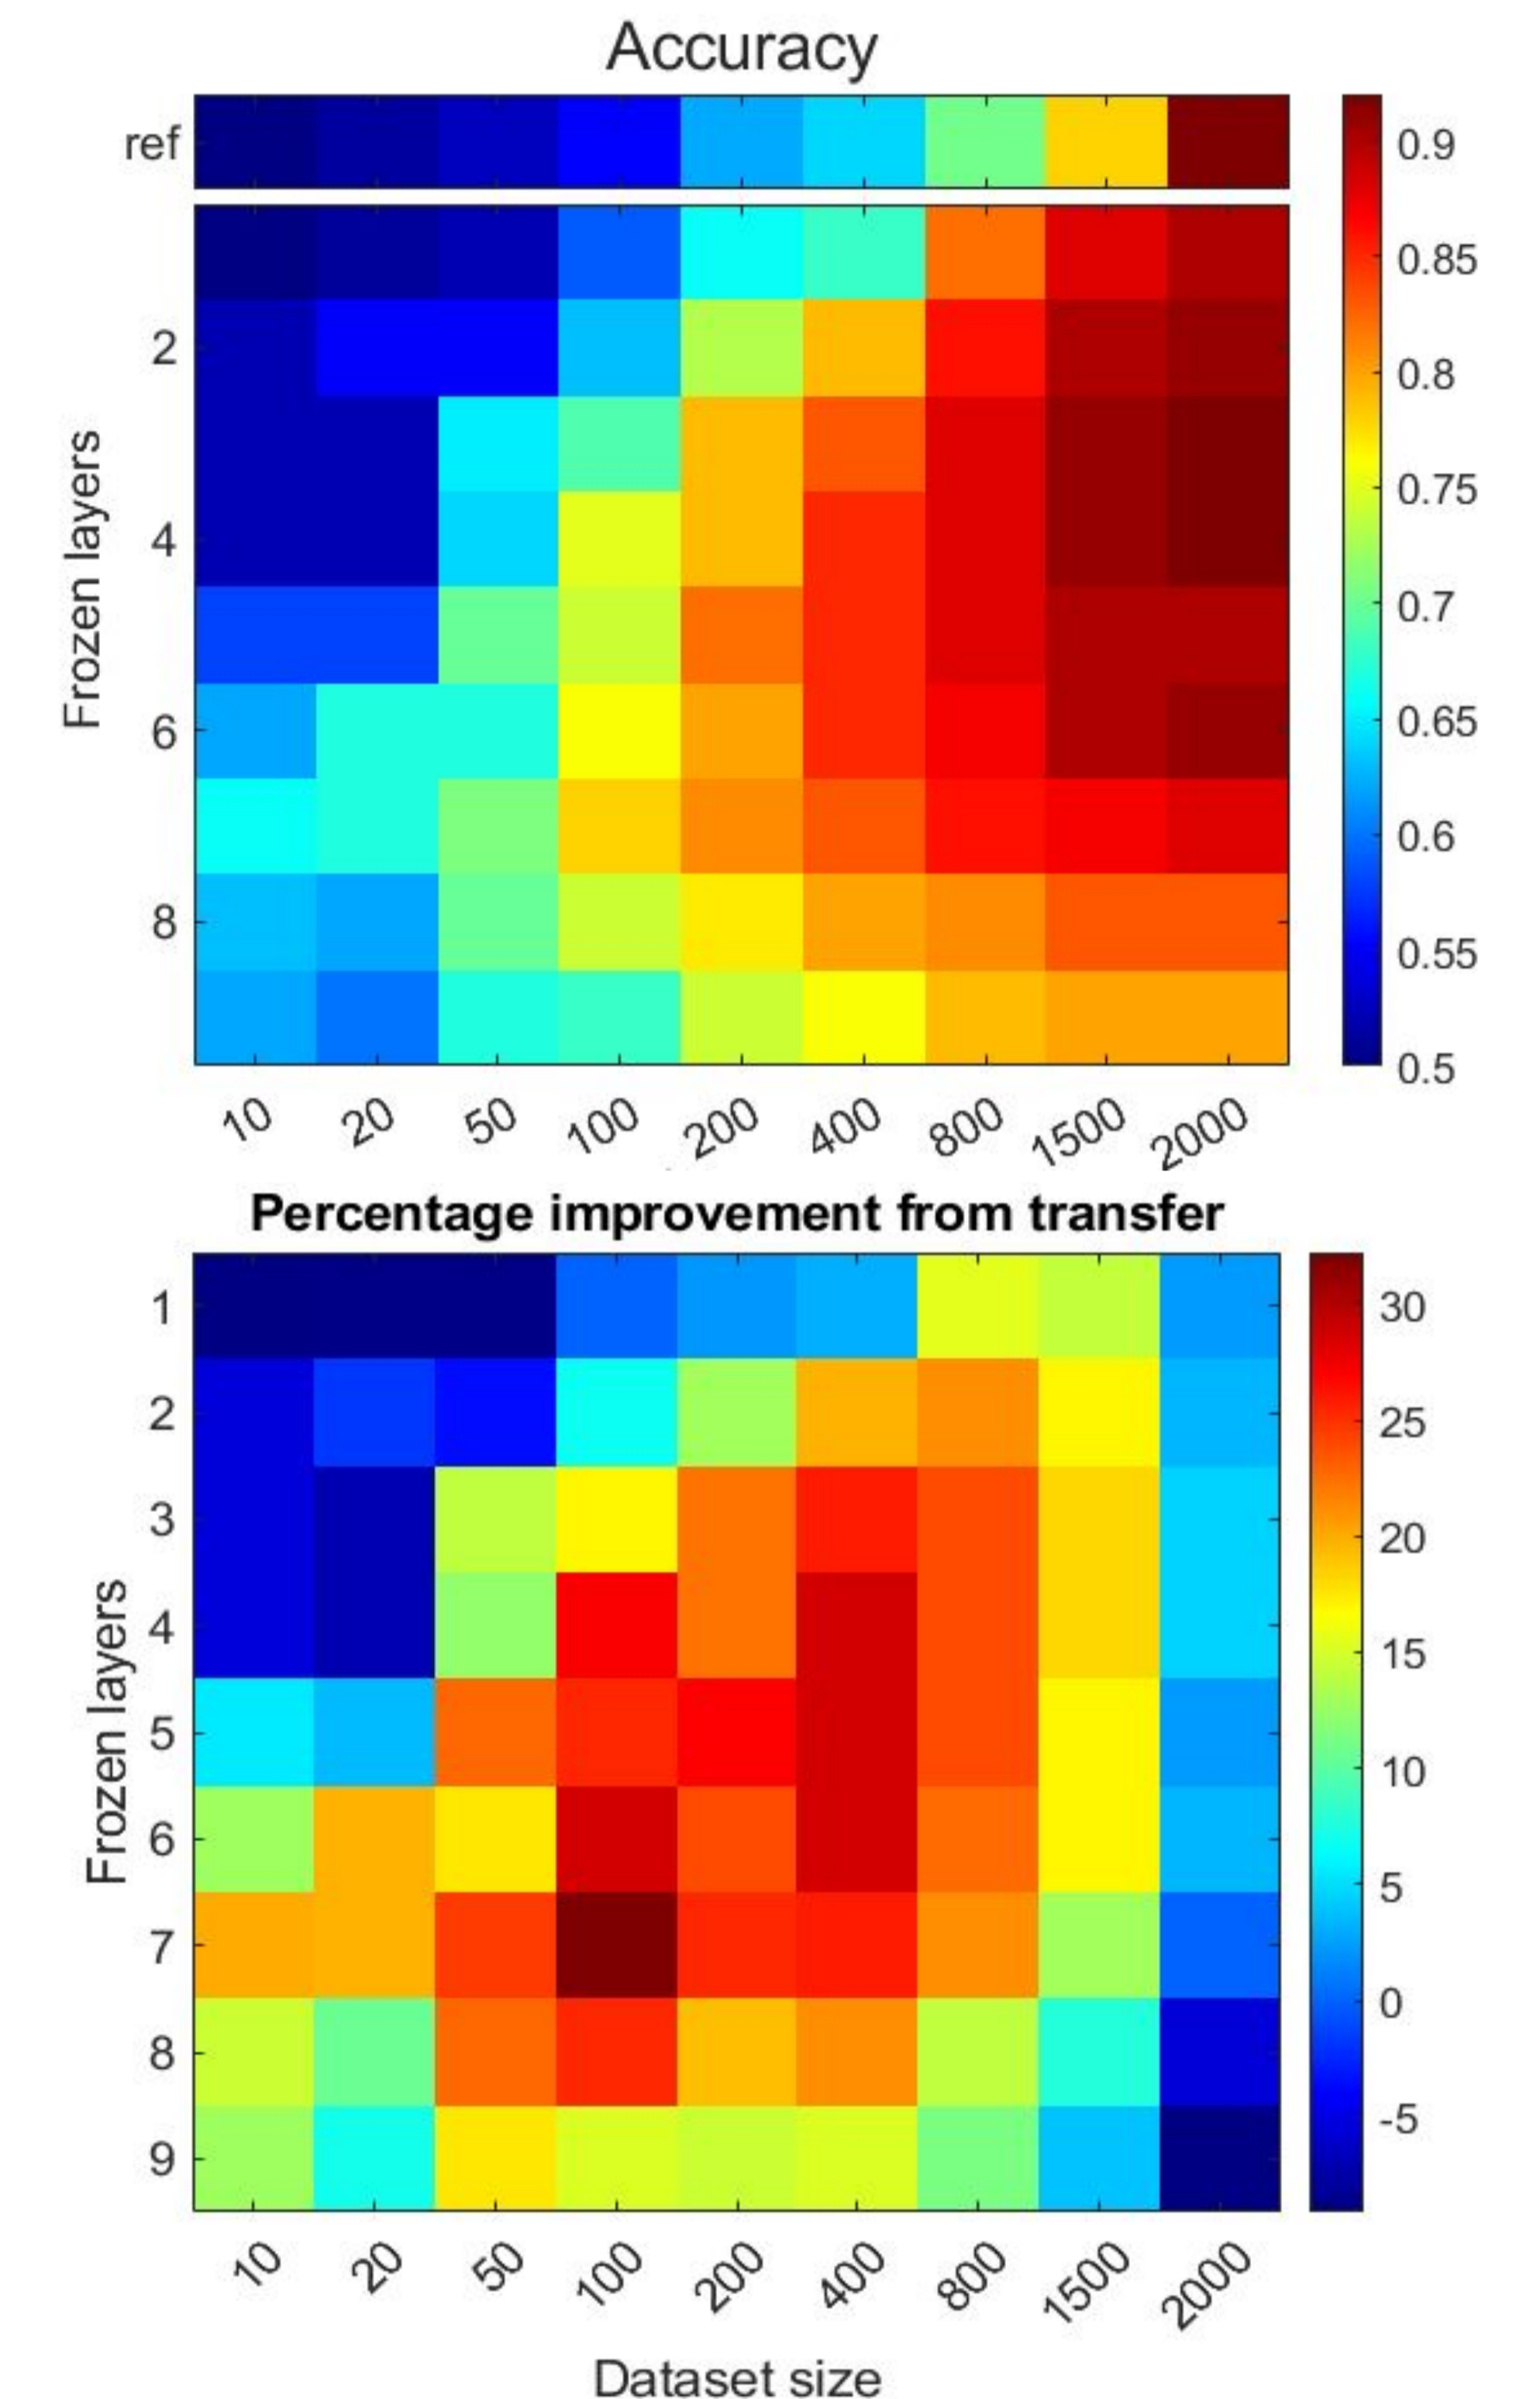


Fig 3: shows experimental results when using TL between domains within the Cifar-10 dataset

Conclusion

- TL may indeed improve performance and learning rate. However it might also hamper accuracy if applied carelessly.
- The number of layers that should be frozen depends on the amount of data points.
- The result does not show to what extent this result is generalisable to other target and source domains.

A project by:
Måns Bengtson
 mabe0107@student.uu.se
Oskar Anderhagen Holmes
 osho2304@student.uu.se
Oscar Holmqvist
 osho5560@student.uu.se

Supervisor:
Jimmy Pettersson
 jimmy.pettersson@hpcsweden.se