Transfer learning with deep convolutional neural networks for classifying cellular morphological changes

UPPSALA UNIVERSITET

Alexander Kensert[a,*], Philip J Harrison[a] and Ola Spjuth[a]

[a] Department of Pharmaceutical Biosciences
* alexander.kensert@gmail.com

Overview

Classical image analysis for cellular phenotyping requires several non-trivial and independent analysis steps. Deep learning through **convolutional neural networks (CNNs)** (Figure 1) has emerged as a compelling alternative to replace these traditional workflows with a single network architecture.

Transfer learning - the transfer of knowledge between tasks - is often beneficial when a limited amount of annotated data is available. Furthermore, CNNs trained on biomedical images, captured under specific experimental condition and imaging setups, can have poor generalizability. To overcome these limitation large annotated datasets, like **ImageNet**, can be used to pre-train state-of-the-art CNNs.

In this study we applied pre-trained CNNs to predict cell mechanisms of action (MoAs) in response to chemical perturbations for two cell profiling datasets from the **Broad Bioimage Benchmark Collection (bbbc)** and obtained higher predictive accuracy than previously reported, between 95 and 97%.

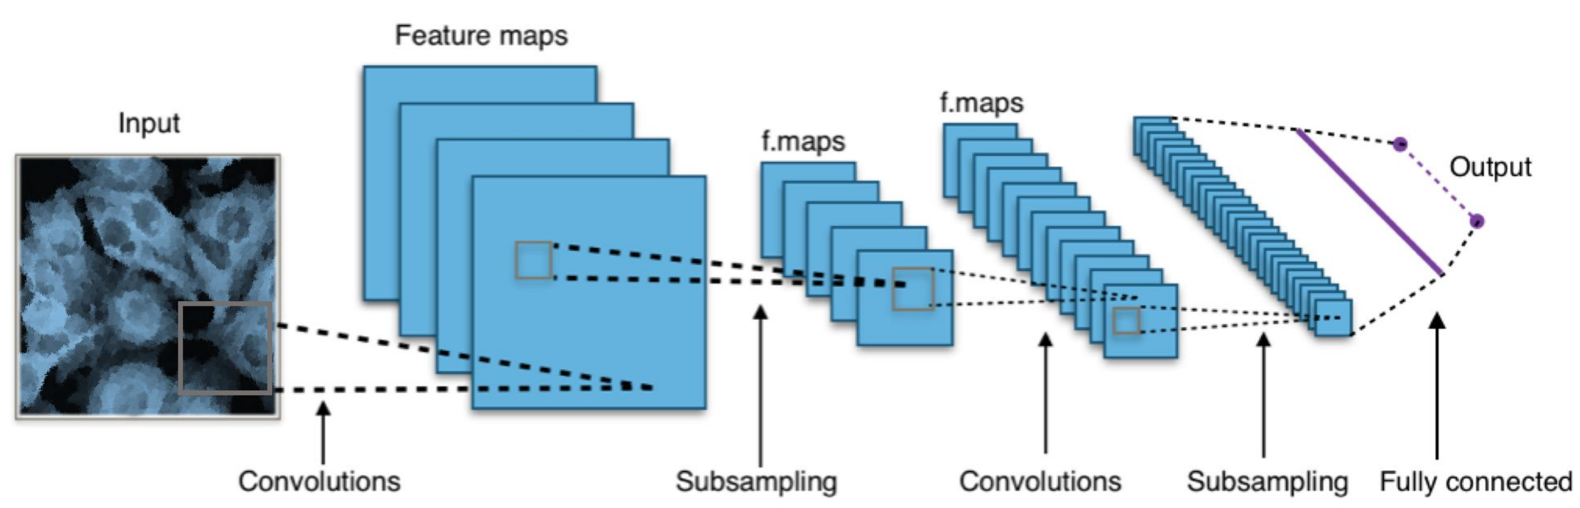

Figure 1. A typical illustration of a CNN. With today's computers, much deeper networks are applied with great predictive power on image classification tasks. Image modified from Wikimedia Commons File:Typical cnn.png

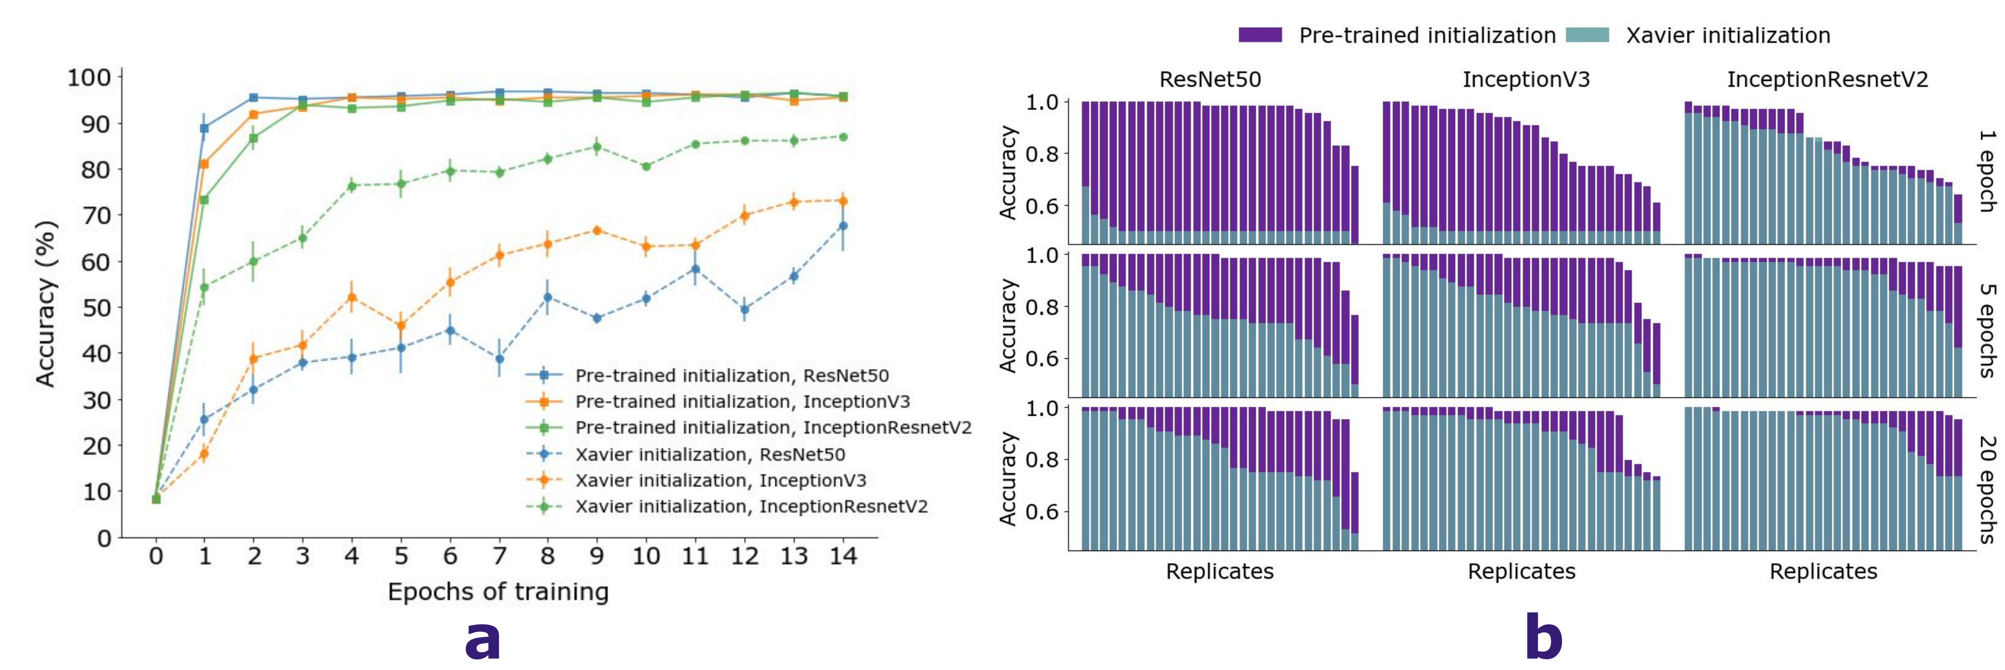

Figure 4. A comparison of test set accuracy between pre-trained applications and Xavier initialized applications (not pre-trained) of the same architectures and the same hyperparameter settings. The plots illustrate how pre-training greatly improves learning.

Material & Methods

Datasets

The CNNs were used to predict mechanisms of action (MoA) and nucleus translocation (Figure 2), based only on pixel intensities which automatically pass through the network to give the final predictions. We used two different bbbc datasets: bbbc021v1 (**MoA dataset**, predicting 12 class labels) and bbbc014v1 (**translocation dataset**, predicting 2 class labels), to evaluate the models' predictive performance.

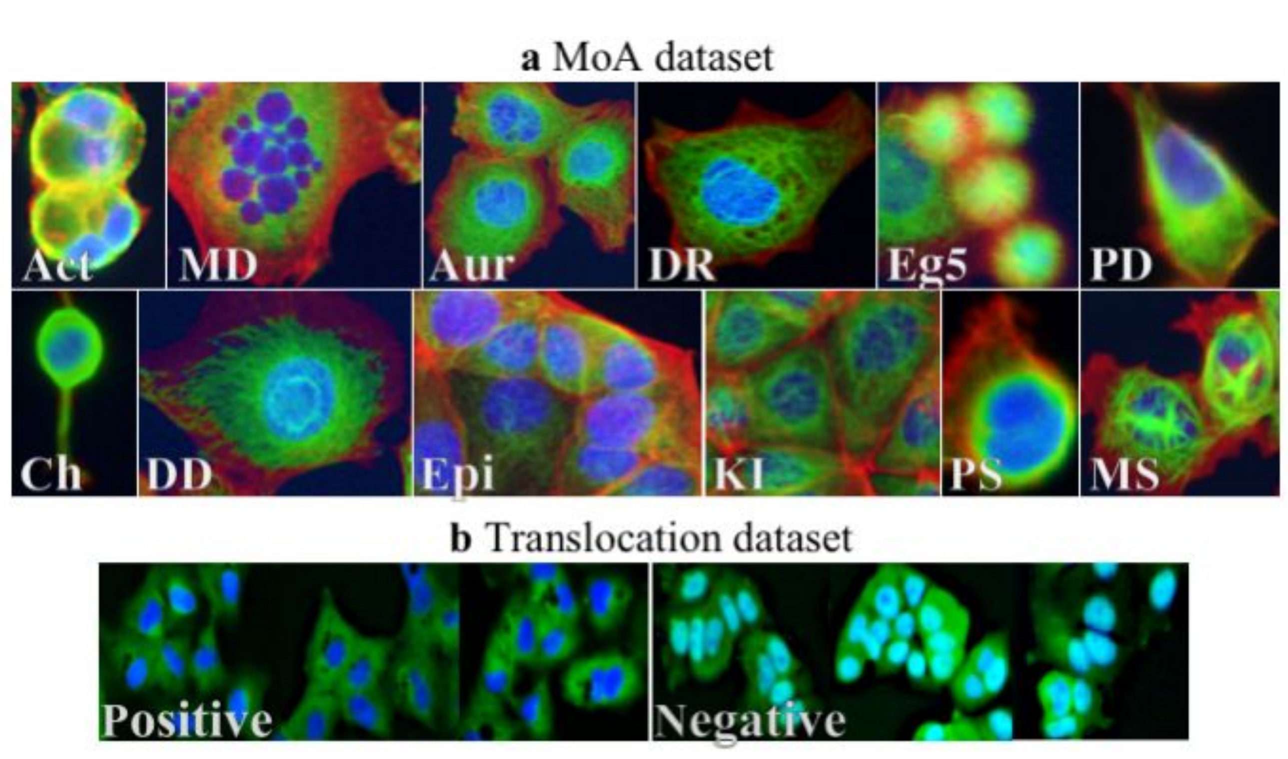

Figure 2.

a illustrates the different mechanisms of actions in the MoA dataset. (see Figure 5 for full names of the abbreviated MoAs).

b illustrates the two different classes (positive and negative) in the translocation dataset. Positive for translocation and Negative for no translocation.

Architectures

Three different architectures were implemented in **TensorFlow** via **Keras**: **Resnet50**, **InceptionV3** and **InceptionResnetV2**. They were all pre-trained on the **ImageNet** dataset, containing 13 million natural images.

Resnet50: 50 layers deep. Includes residual mappings to enable the fitting of deeper and thus more discriminating networks than would otherwise be possible.

InceptionResnetV2: 245 layers deep. Combines both inception blocks and residual mappings.

InceptionV3: 95 layers deep. It is not always certain what filter sizes to use for the convolutions (Figure 3), to overcome this Inception architectures include multiple filter sizes for the network to pick from given the data at hand.

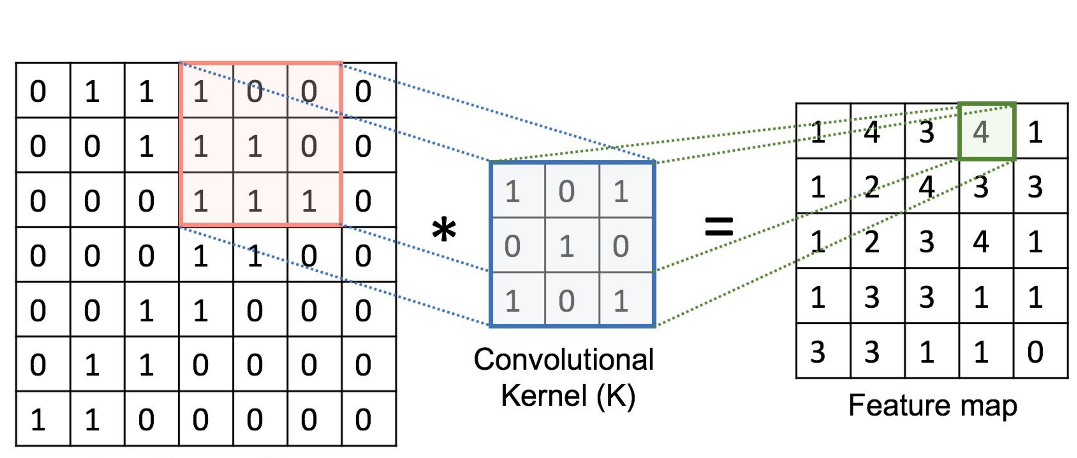

Figure 3. An example of an input image I convolved with a filter $K_{3\times3}$ with weights of zeros and ones to encode a representation (feature map). The receptive field is highlighted in pink and the corresponding output value for the position is marked in green. Figure courtesy of Anindya Gupta.

Results

We illustrate the prediction accuracies on the **MoA dataset** across epochs of training (Figure 4**a**) and as confusion matrices (Figure 5). **ResNet50**, **InceptionV3** and **InceptionResnetV2** attained mean accuracies of 97%, 97% and 95% respectively – thus reaching greater accuracy than any model yet reported based on this dataset. However, although our models correctly predicted the MoA for these treatments, there were still high **uncertainties** in several of the predictions.

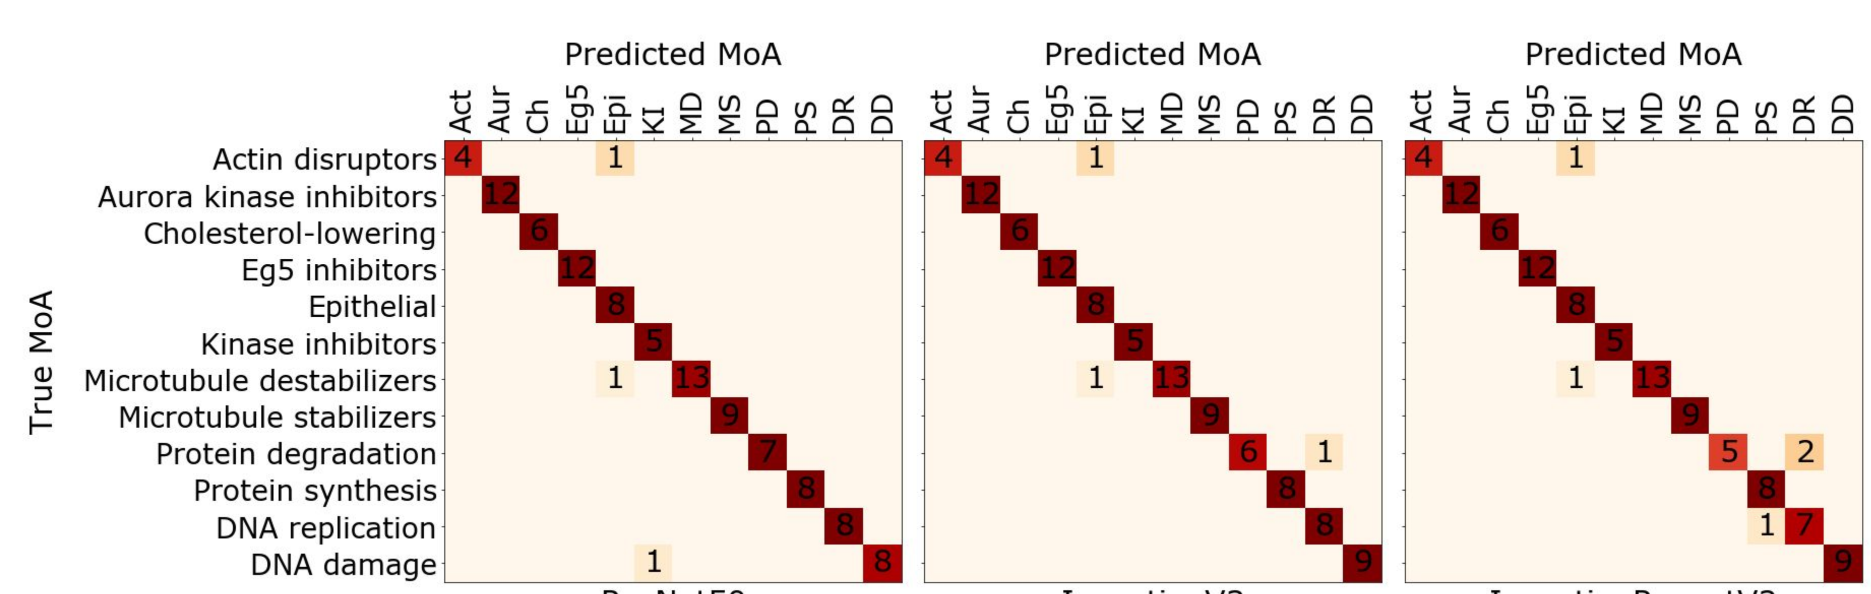

Figure 5. Confusion matrices for hard predictions of compound-concentration pairs. Zeros are excluded for better visualization.

On the **translocation dataset** the three models attained accuracies of up to 100% after just single epochs of training (Figure 4**b**). The quick learning is arguably a strong indication of transferability of the pre-trained parameters.

Conclusions & Future work

Transfer learning allows the fitting of deeper networks based on fewer task-specific annotated images. It also gives faster convergence (i.e. fewer training epochs are required) and improved classification performance and generalizability.

As mentioned earlier, there were high **uncertainties** in many of our predictions. Formally quantifying and accounting for this uncertainty is of significant interest. In future work we plan to explore various means of doing this (including **conformal prediction** and **Bayesian methods**) to extend the work presented here.

pharmb.io
HASTE
SWEDISH FOUNDATION for STRATEGIC RESEARCH