Implementing alternative network architectures for protein-ligand CNNs

Lillian Turner,[1] Pulkit Mitttal,[1,2] and David Ryan Koes[3]

[1]Department of Computer Science, [2]Department of Biology, [3]Department of Systems and Computational Biology

University of Pittsburgh

Background

Drug discovery can be both expensive and time-consuming. Using machine learning, we hope to reduce these costs by using computers to predict protein-ligand binding and quickly identify useful drugs.

Computers can learn spatial relationships using convolutional neural nets, or CNNs. CNNs take in data and pass it through a series of nodes to generate a prediction from the input. These nodes compute non-linear functions of regions of the input and are organized in layers that maintain spatial relationships. Successive layers can learn higher-order features of the input to successfully recognize complex relationships.

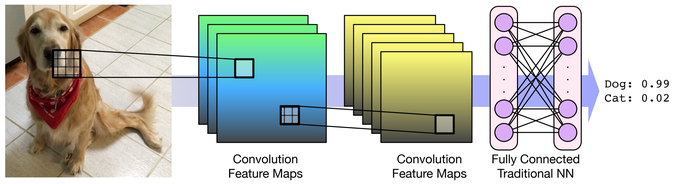

During training, the network is shown labeled data. If the prediction does not match the label, the weights on different components of the network are adjusted using stochastic gradient descent. After many iterations the network can make useful predictions about test data that was not part of the training set.

In this experiment, we look at three different variants of CNNs – densely connected neural networks (DenseNets), residual networks (ResNets), and GoogLeNet. These types of networks have shown strong performance in image recognition. In this experiment, we evaluate their performance in predicting protein-ligand binding poses and affinities.

Methods

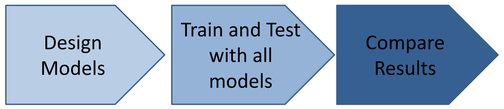

The Caffe and gnina (http://github.com/gnina) frameworks are used to train models. More than 250,000 docked structures generated from the PDBbind 2016 refined set were used as input. Models are trained both for pose scoring, distinguishing low-RMSD from high-RMSD poses, and affinity prediction. The dense networks were evaluated multiple times using different random seeds and clustered cross-validation train/test splits and were trained for 100k iterations. Less dense models were evaluated a single time using randomized (non-clustered) train/test splits and were trained for 50k iterations.

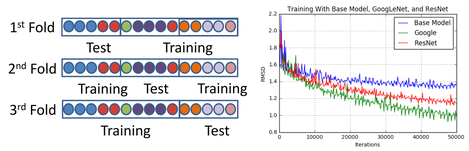

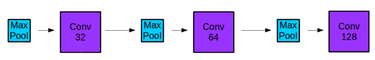

Base model

Densely connected neural nets∗

(and non-densely connected counterparts)

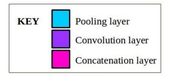

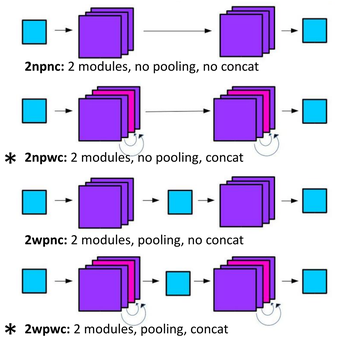

2npnc: 2 modules, no pooling, no concat

∗ **2npwc:** 2 modules, no pooling, concat

2wpnc: 2 modules, pooling, no concat

∗ **2wpwc:** 2 modules, pooling, concat

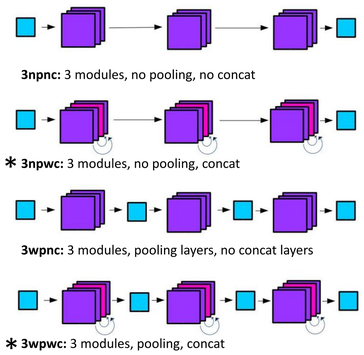

3npnc: 3 modules, no pooling, no concat

∗ **3npwc:** 3 modules, no pooling, concat

3wpnc: 3 modules, pooling layers, no concat layers

∗ **3wpwc:** 3 modules, pooling, concat

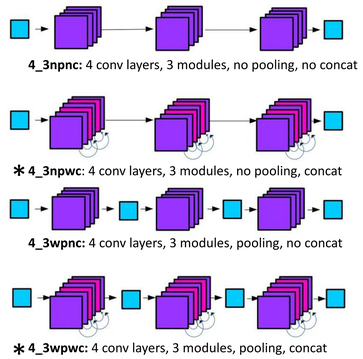

4_3npnc: 4 conv layers, 3 modules, no pooling, no concat

∗ **4_3npwc:** 4 conv layers, 3 modules, no pooling, concat

4_3wpnc: 4 conv layers, 3 modules, pooling, no concat

∗ **4_3wpwc:** 4 conv layers, 3 modules, pooling, concat

Less densely connected neural nets

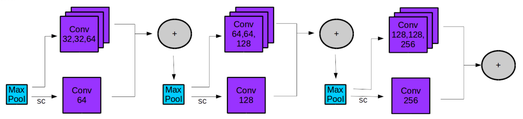

ResNet

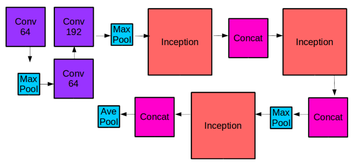

GoogLeNet

Results

Densely connected neural nets

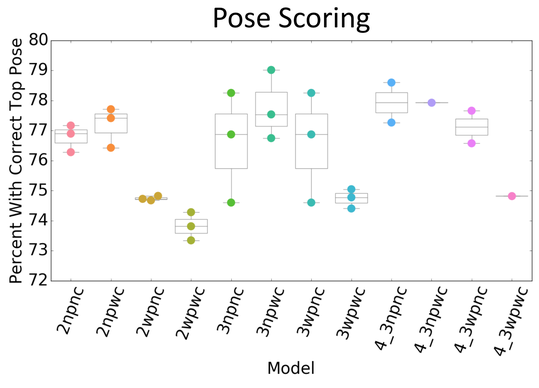

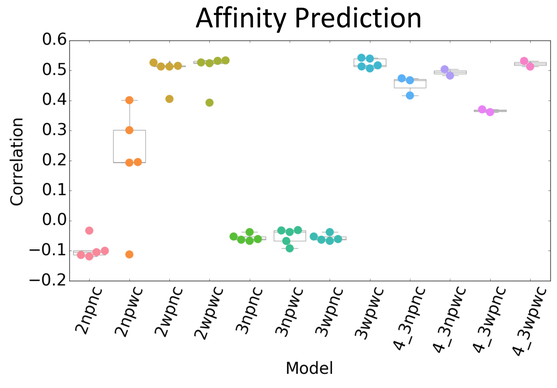

Less densely connected neural nets

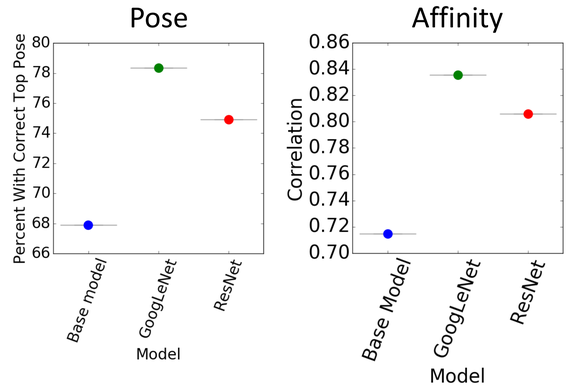

Combined test performance on the pose scoring and affinity prediction tasks is shown. Pose scoring is evaluated by the percentage of targets where a low RMSD (<2Å) pose is top ranked. Affinity prediction is evaluated using Pearson R correlation. Note that the performance of the DenseNet models and the ResNet and GoogLeNet models cannot be compared as they were not evaluated using the same procedure.

Discussion

Our results show that DenseNets outperform their non-densely connected counterparts at the pose prediction task when there are no pooling layers between modules. DenseNets uniformly improve the affinity prediction task. Interestingly, the performance on the pose scoring and affinity prediction tasks tended to vary inversely.

The DenseNets model that did the best overall was the model with the most parameters, 4_3npwc.

Out of the less densely connected neural nets, GoogLeNet performed best. ResNet also performed substantially better than the base model. Unlike with DenseNets, the trends between pose scoring and affinity prediction were the same. These results will be verified with more trials using a more robust evaluation framework.

Overall, we find denser, more connected networks are capable of learning better performing neural network models.

This work is supported by R01GM108340 from the National Institute of General Medical Sciences.