

# Residual Attention Networks for Image Classification

Thaminda Edirisooriya (tediris@stanford.edu)

## Introduction

The application of attention mechanisms to image classification is a newer idea among the many models that exist for this task

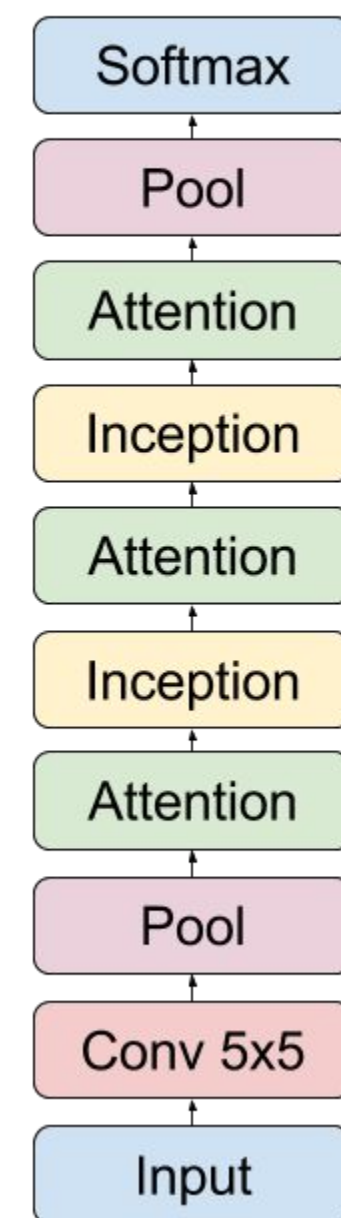
## Task: Imagenet-200

For this task, we seek to correctly classify each image as one of 200 classes. Each image is 64x64 RGB.

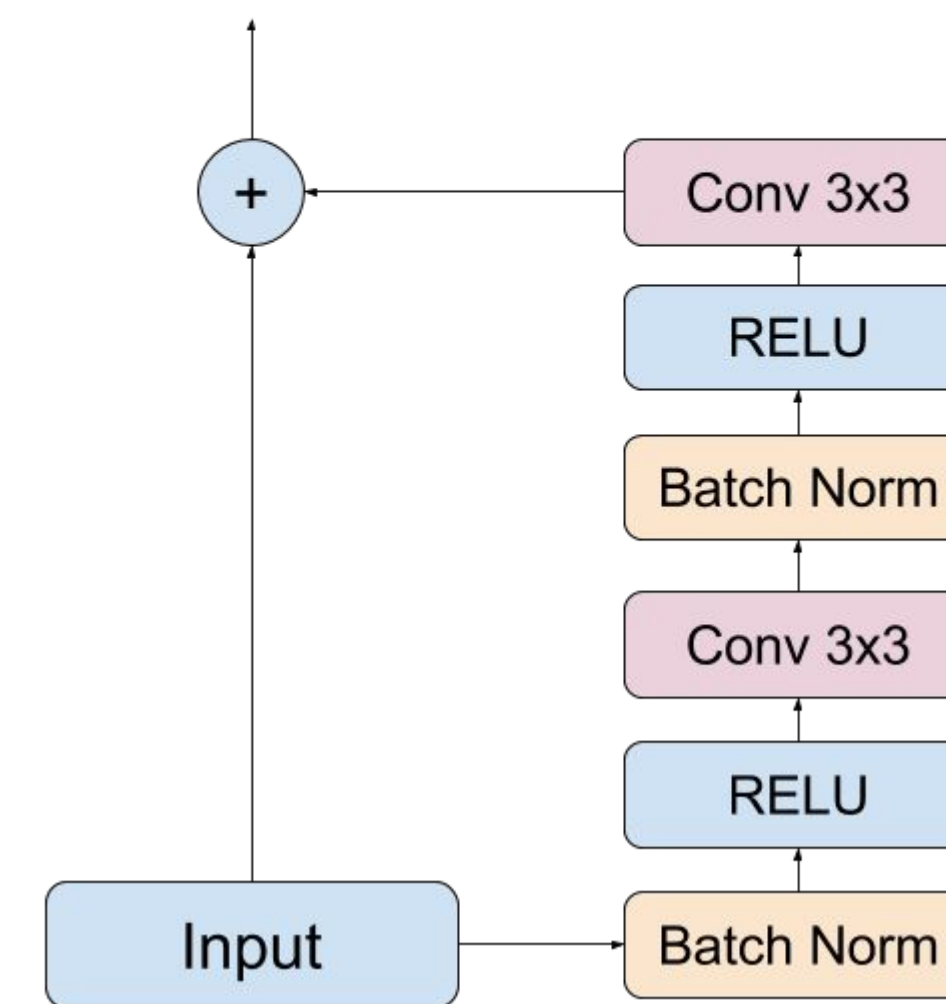
## Experiments

Besides developing a network structure that can use attention effectively, I've been experimenting with different residual modules (ResNet, Inception Residual V2)

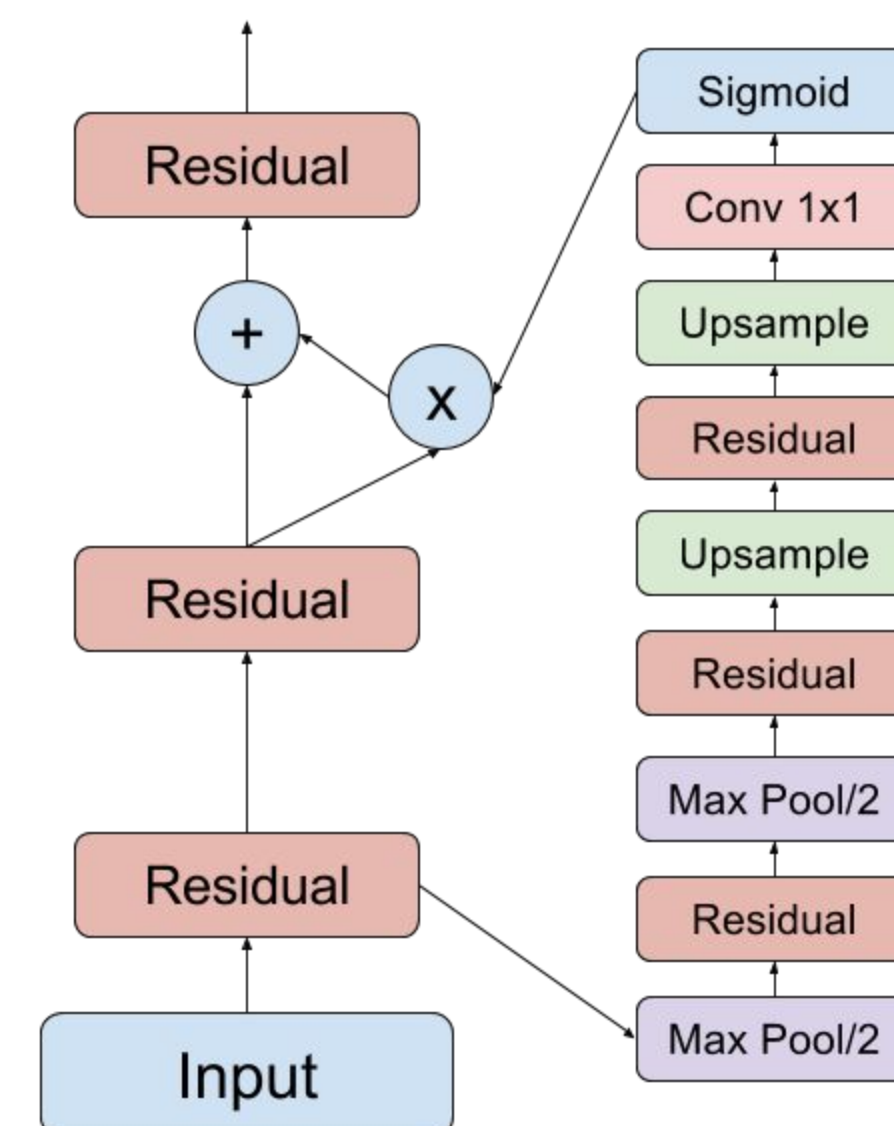
### Network Architecture



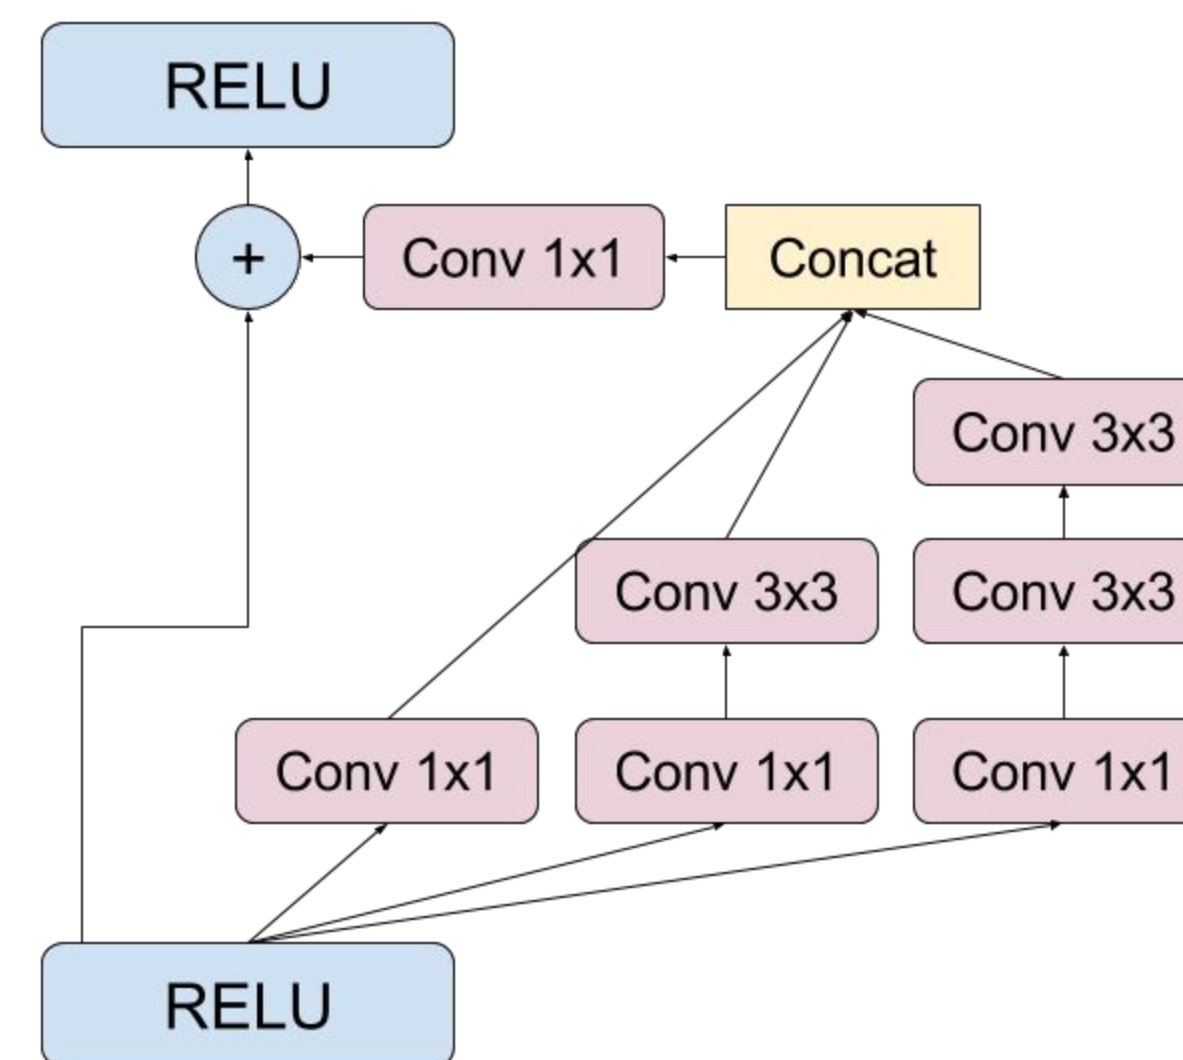
### Residual Module with Pre-Activation




### Attention Module




### Residual Inception Module



Validation Accuracy: Simple Residual vs Inception



Validation Loss with Simple Residual vs. Inception Residual Attention



## Conclusion

From the results, we see that the network appears to train better when using the simpler residual modules than using many inception modules. I plan to focus on this residual module choice, to tune hyperparameters and create a performant system

### References

Wang, F. & Jiang, M & Qian, C: Residual Attention Network for Image Classification

Szegedy, C & Ioffe, S & Vanhoucke, V & Alemi, A: Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning