MagmaDNN: Applications in Microscopy

Students: Sedrick Keh (HKUST), Daniel Nichols (UTK), Kam Fai Chan (CUHK)
Mentors: Junqi Yin, Kwai Wong

Computational Microscopy: Space Group Classification:

Every material can be associated with a particular symmetric space group. Though it may seem that materials can have infinitely many possible configurations, researchers have discovered that there are in fact only 230 possible symmetric groups.

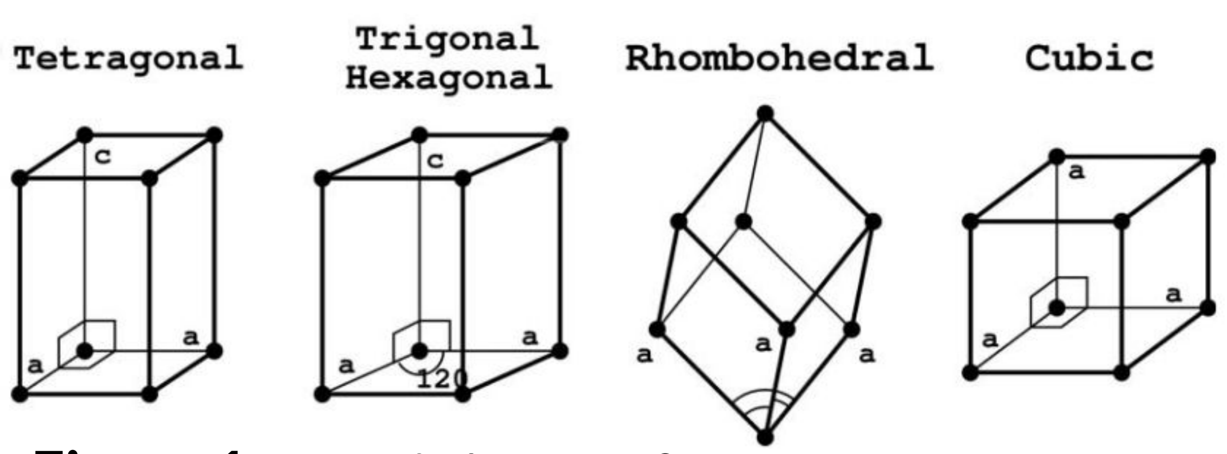

Figure 1: *Broad classes of space groups*

Task:

The main objective is to apply deep learning to space group classification. The dataset* consists of (3 x 512 x 512) image intensities of convergent beam electron diffraction (CBED) patterns, where each channel represents a different material projection.

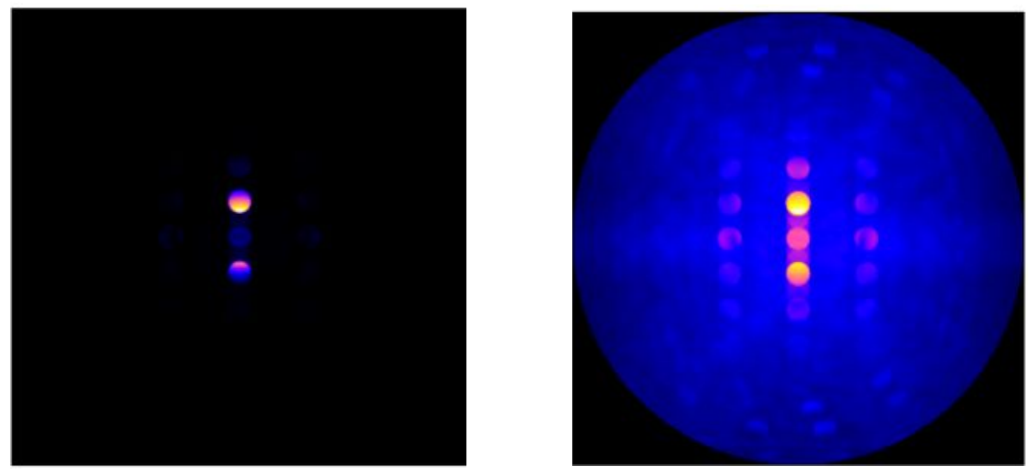

Figure 2: *Visualized CBED image intensities, unscaled (left) and scaled (right)*

Challenges:

The main challenge intrinsic to the data is its large size, meaning training time may be long. Thus, training a model for this task would likely require data parallelism or model parallelism (likely both).

*This task is taken from the Smoky Mountain Data Challenge 2019

How can MagmaDNN be used in this task?

MagmaDNN is a modularized deep learning framework that is optimized for parallel computation and distributed training on GPUs.

Accelerated GPU Computation
MagmaDNN is built around the MAGMA library for optimized linear algebra routines and the CuDNN library for optimized DNN-related algorithms on the GPU.

Dynamic Memory Manager
As memory operations (e.g. copying from the CPU to GPU) may have large overhead costs, MagmaDNN defines its own memory manager class, leading to large speedups in model training.

ResNet-50:

This is an application that combines most of MagmaDNN's core features. It is a full-scale deep architecture that stacks multiple convolutional layers, together with common regularization and downsampling layers like pooling and dropout.

A residual network (ResNet) [1] supports very deep convolutional networks by adding a residual shortcut layer that prevents the gradient from exponentially exploding or shrinking. ResNet-50 is a 50 layer residual network with additional bottleneck layers that use 1 by 1 projection layers.

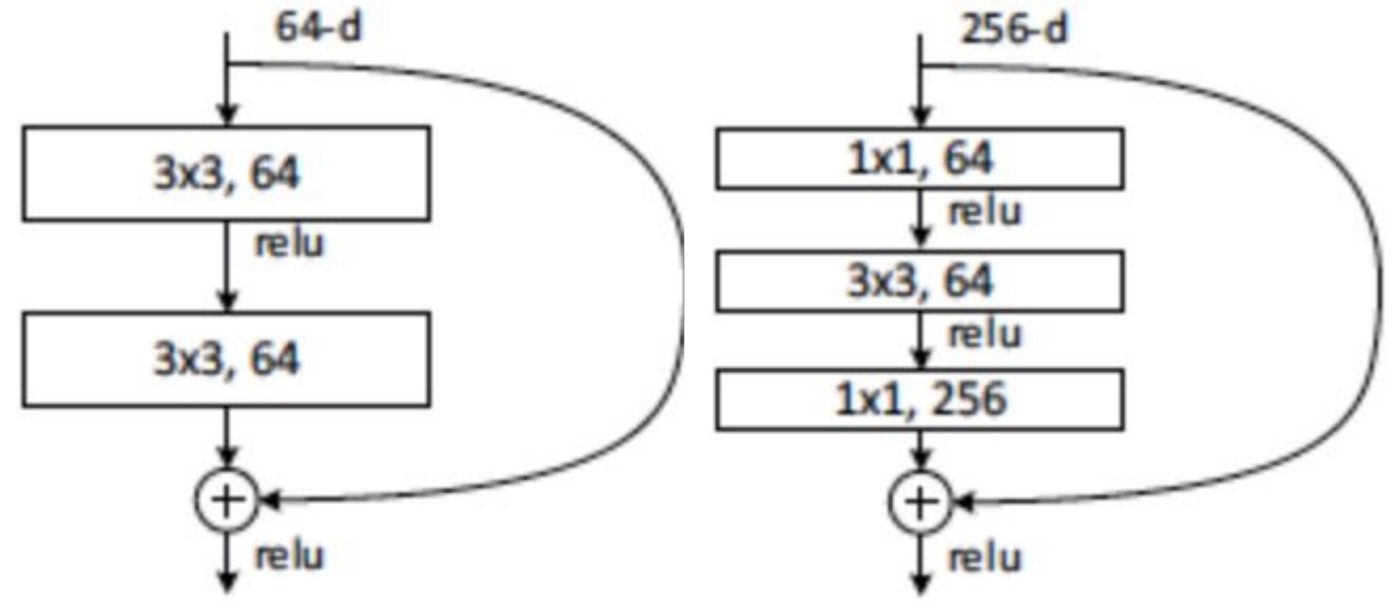

Figure 3: *ResNet skip connections (left), skip connections with bottleneck (right)*

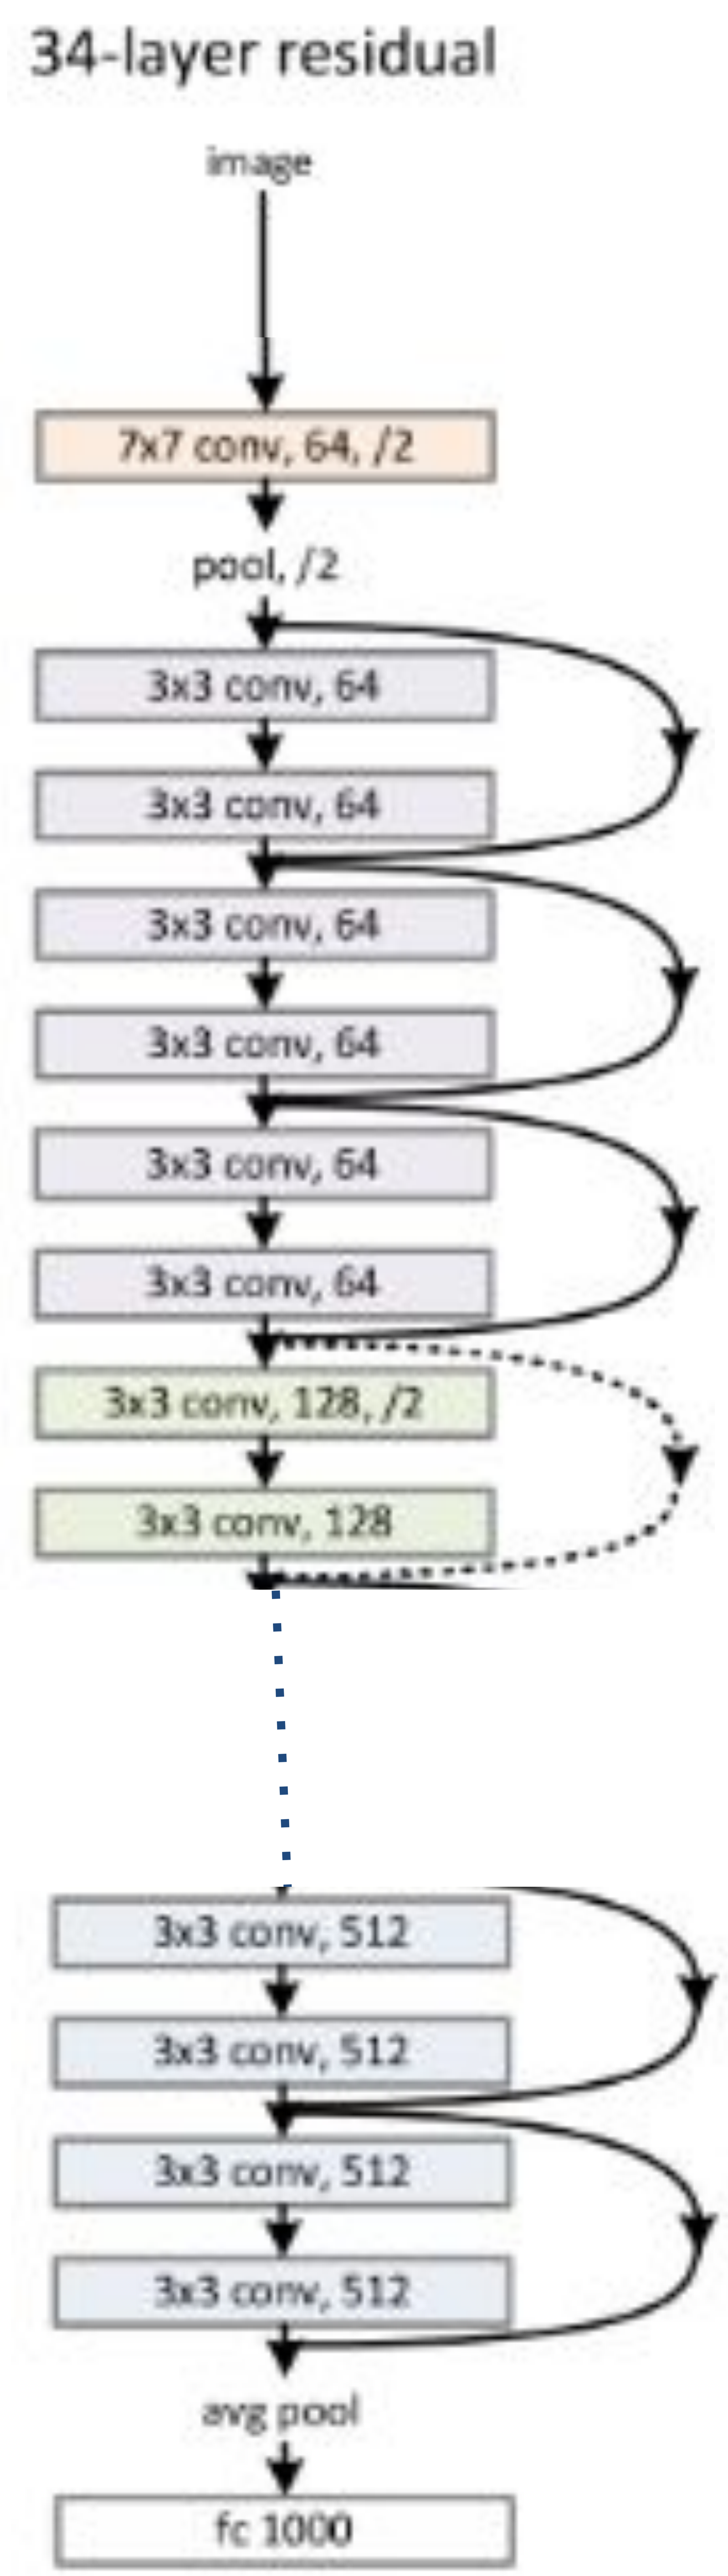

Figure 4: *ResNet Architecture*

[1] K. He, et al. *Deep Residual Learning for Image Recognition.*

Results and Analysis

MagmaDNN scales well with increased data. The table below shows how the time scales as the number of iterations increase.

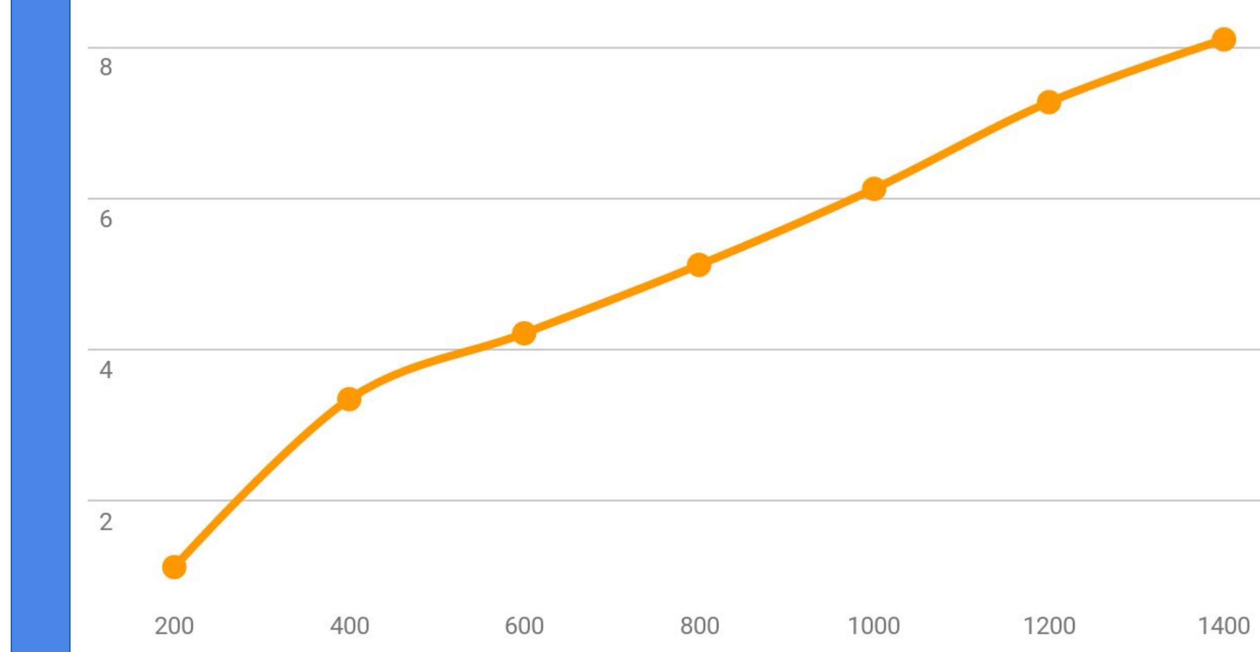

Figure 5: *Time vs Iterations Graph*

At the moment, ResNet-50 has still not been used with the CBED image dataset. Rather, we begin with a simple convolutional layer with two layers. With a stochastic gradient descent optimizer, a constant learning rate of 0.0001 and a batch size of 32, the model was able to reach 0.16 accuracy after processing 4500 images. This is expected to increase as larger models are used.

Future Work

The next step is to apply the ResNet-50 model to the CBED image dataset. This can also be expanded into ResNet-152 or even ResNet-1k. Additionally, more features (such as batch normalization) can be further added to MagmaDNN to improve the quality of the model.

Acknowledgements

This work was done under the RECSEM REU program at the University of Tennessee under the supervision of Dr. Kwai Wong, with support from the National Science Foundation, Joint Institute of Computational Sciences (UT-ORNL), Extreme Science and Engineering Discovery Environment (XSEDE), BP HPC team.