Making Deep Learning Models Memory Efficient

R&D SHOWCASE

INTERNATIONAL INSTITUTE OF INFORMATION TECHNOLOGY HYDERABAD

MOTIVATION

- Deep Learning Models needs to run on small devices.
- Model needs to process video at 30 fps.

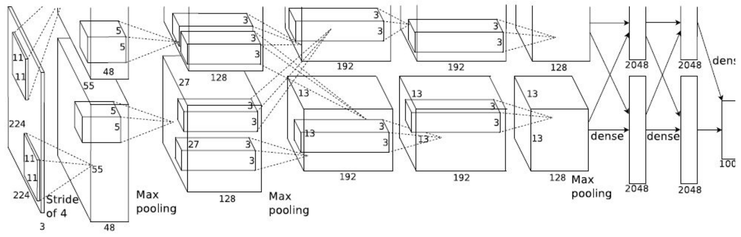

- But current models Eg.-AlexNet has 60 million parameters (~240 MB on disk) and performs 1.5 billion single precision operations to classify one image i.e in a forward pass.)

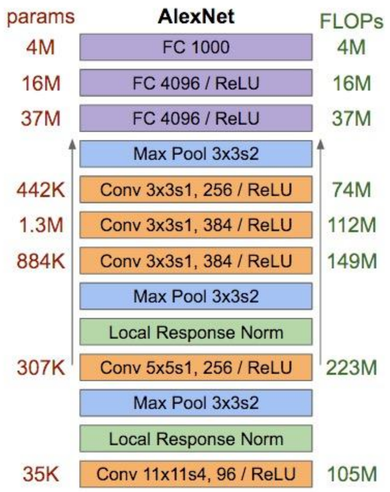

METHODS

- Pruning

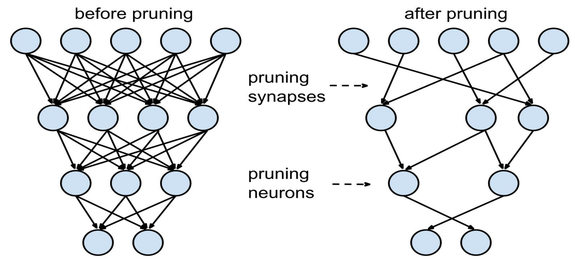

- Quantization

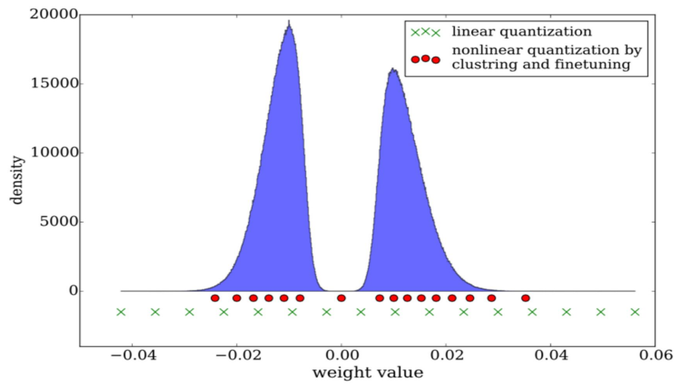

- Pruning and Quantization while Training

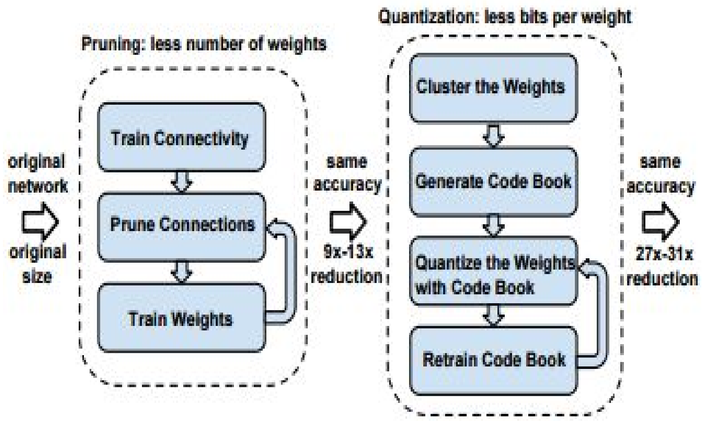

RESULTS

- Quantization results for Image Classification with Imagenet

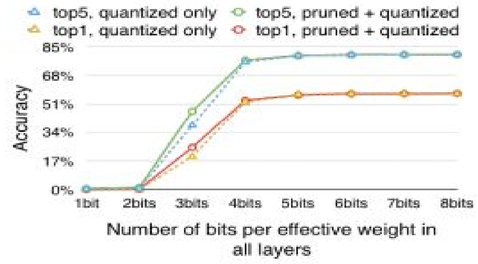

- For Image Retrieval Dataset

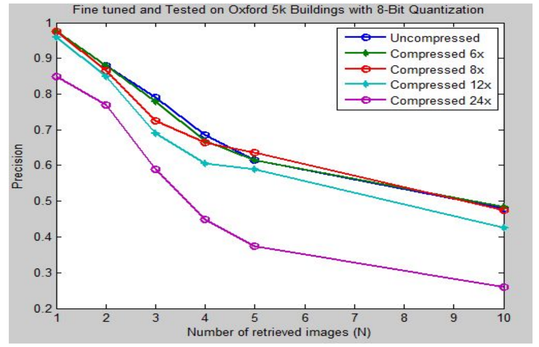

Qualitative results for Image Search.
The 1st, 2nd and 3rd rows are retrieved images with Ground Truth,Uncompressed Network and 12X Compressed Network.

Soham Saha, Girish Varma, C.V.Jawahar, Compressing Deep Neural Networks for Recognizing Places, ACPR, 2017

Contact: soham.saha@research.iiit.ac.in, girish.varma@iiit.ac.in

Machine Learning Lab, KCIS