Half the Measurements, Twice the Speed

Accelerating Deep RL With Compressed Sensing

Michael Teti, Dr. William Hahn, and Dr. Elan Barenholtz
Florida Atlantic University
Max Planck Florida Institute for Neuroscience

NICE 2019

DNNs Are Really Good At a Lot of Things

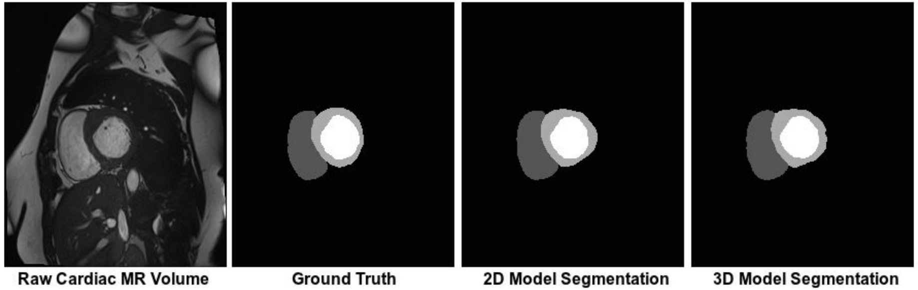

"If a typical person can do a mental task with less than one second of thought, we can probably automate it using AI either now or in the near future."

-Andrew Ng

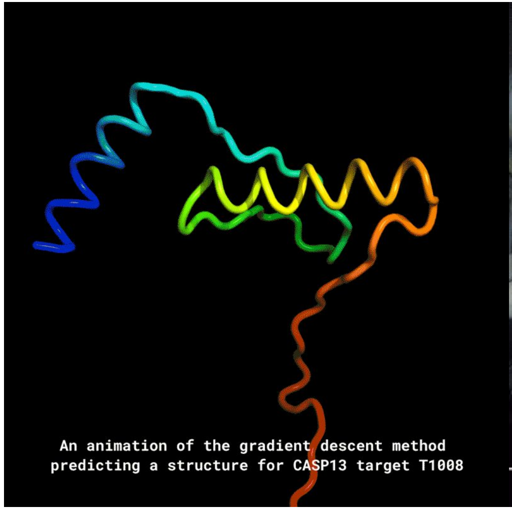

They Also Have Their Disadvantages

- Require huge, diverse, and labeled datasets to avoid overfitting
  - Transfer learning, data augmentation, dropout, and regularization often help
- Long training times
  - Transfer learning, normalization
- Millions of tunable weights and huge, complicated architectures
  - Transfer learning (again),, weight compression, GPUs, global average pooling
- Various architectures, and hyper-parameters that are often problem-specific
  - hyperparameter tuning, trial-and-error, https://stackexchange.com/

They Also Have Their Disadvantages

- Require huge, diverse, and labeled datasets to avoid overfitting
    - Transfer learning, data augmentation, dropout, and regularization often help
- Long training times
    - Transfer learning, normalization
- Millions of tunable weights and huge, complicated architectures
    - Transfer learning (again),, weight compression, GPUs, global average pooling
- Various architectures, and hyper-parameters that are often problem-specific
    - hyperparameter tuning, trial-and-error, https://stackexchange.com/

They Also Have Their Disadvantages

- Require huge, diverse, and labeled datasets to avoid overfitting
  - Transfer learning, data augmentation, dropout, and regularization often help
- Long training times
  - Transfer learning, normalization
- Millions of tunable weights and huge, complicated architectures
  - Transfer learning (again),, weight compression, GPUs, global average pooling
- Various architectures, and hyper-parameters that are often problem-specific
  - hyperparameter tuning, trial-and-error, https://stackexchange.com/

They Also Have Their Disadvantages

- Require huge, diverse, and labeled datasets to avoid overfitting
    - Transfer learning, data augmentation, dropout, and regularization often help
- Long training times
    - Transfer learning, normalization
- Millions of tunable weights and huge, complicated architectures
    - Transfer learning (again),, weight compression, GPUs, global average pooling
- Various architectures, and hyper-parameters that are often problem-specific
    - hyperparameter tuning, trial-and-error, https://stackexchange.com/

K.I.S.S. (Keep It Simple, Stupid)

- Current trend of bigger, more elaborate networks is not sustainable
  - edge technologies, power consumption, AGI
- Why train a larger network only to prune it later upon deployment?

Inspiration from signal processing

- Compressed sensing (Candes et al. 2005) is a signal acquisition *and* compression method that samples well below the Nyquist frequency
- Robust to noise
- Dumb encoder, smart decoder
- Not interested in decoding here

"One can regard the possibility of digital compression as a failure of sensor design. If it is possible to compress measured data, one might argue that too many measurements were taken."

- David Brady

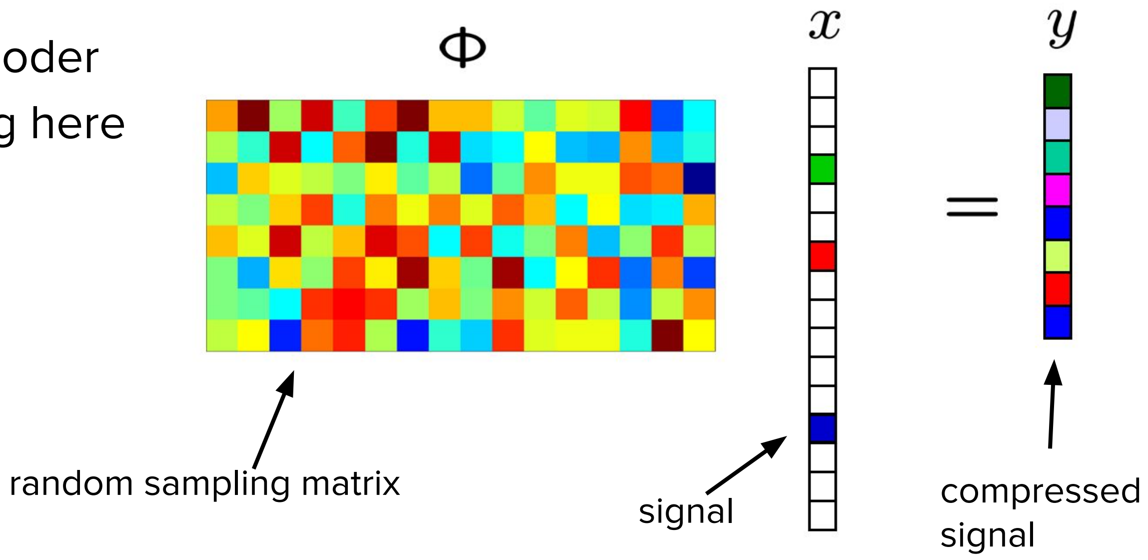

Compressed Learning

- Reconstruction illustrates inherent information is stored in compressed vector
- Reconstruction is cool

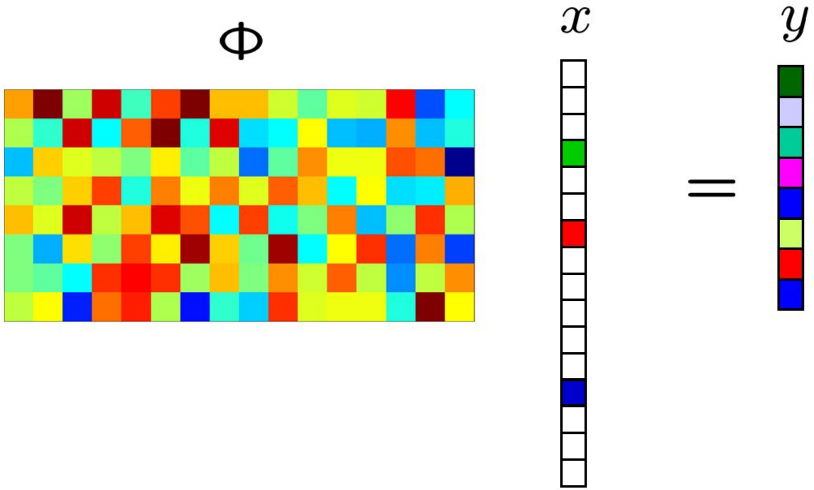

Compressed Learning

- Reconstruction illustrates inherent information is stored in compressed vector
- Reconstruction is cool
- Learning is cooler

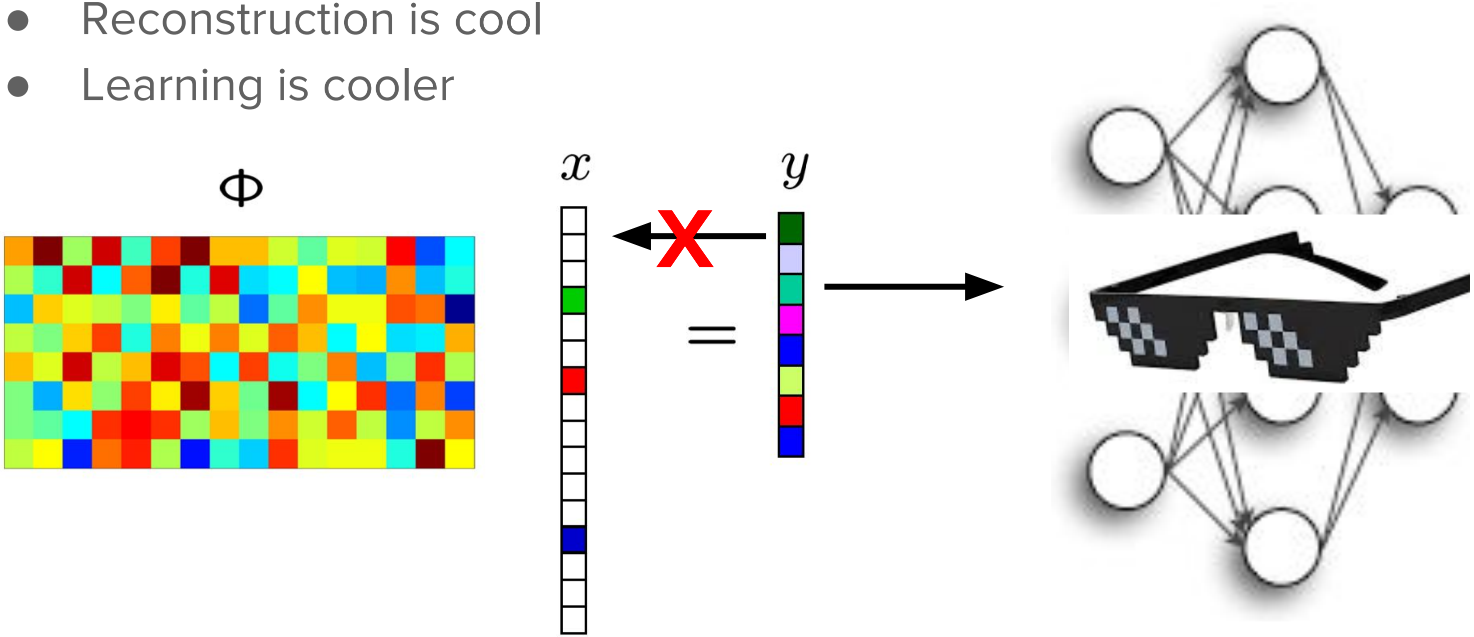

Compressed Sensing and DNNs

- Many approaches concentrate compression effort on inputs
  - Adler et al. (2016)
  - Gueguen et al. (2018)
- Ehrlich et al. (2019) propose a residual framework

Our Method

- We propose randomly projecting the last feature map layer in to a lower-dimensional vector

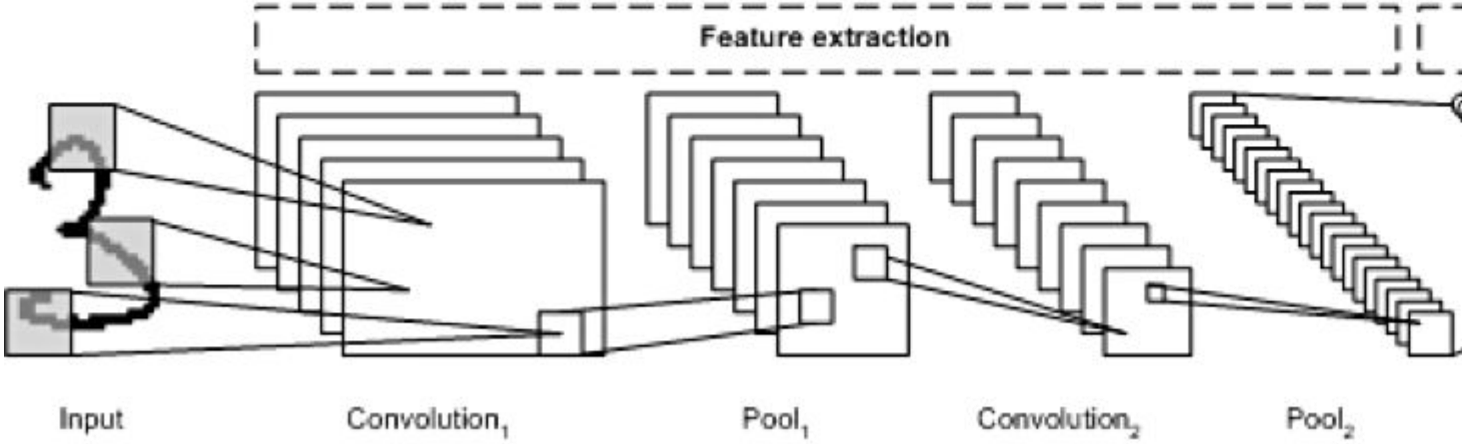

RL Task

- DQN learning to play pong from scratch
- Rewarded or punished by score relative to opponent
- Max score: 21

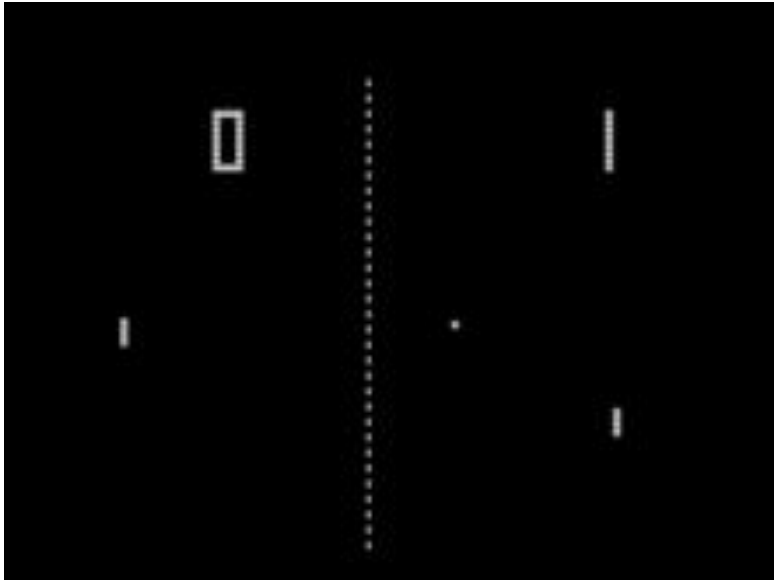

Layer	Filter / Weight Dimension	Output Dimension
Input	4 x 84 x 84	-
Conv1	8 x 8, 32, stride 4	32 x 21 x 21
ReLU	-	32 x 21 x 21
Conv2	4 x 4, 64, stride 2	64 x 10 x 10
ReLU	-	64 x 10 x 10
Conv3	3 x 3, 64, stride 1	64 x 10 x 10
ReLU	-	64 x 10 x 10
Fully-connected1	6400 x 512	512
Output	512 x 6	6

Layer	**Filter / Weight Dimension**	**Output Dimension**
Input	4 x 84 x 84	-
Conv1	8 x 8, 32, stride 4	32 x 21 x 21
ReLU	-	32 x 21 x 21
Conv2	4 x 4, 64, stride 2	64 x 10 x 10
ReLU	-	64 x 10 x 10
Conv3	3 x 3, 64, stride 1	64 x 10 x 10
ReLU	-	64 x 10 x 10
Random Projection	**6400 x # samples**	**# samples**
Fully-connected1	**# samples** x 512	512
Output	512 x 6	6

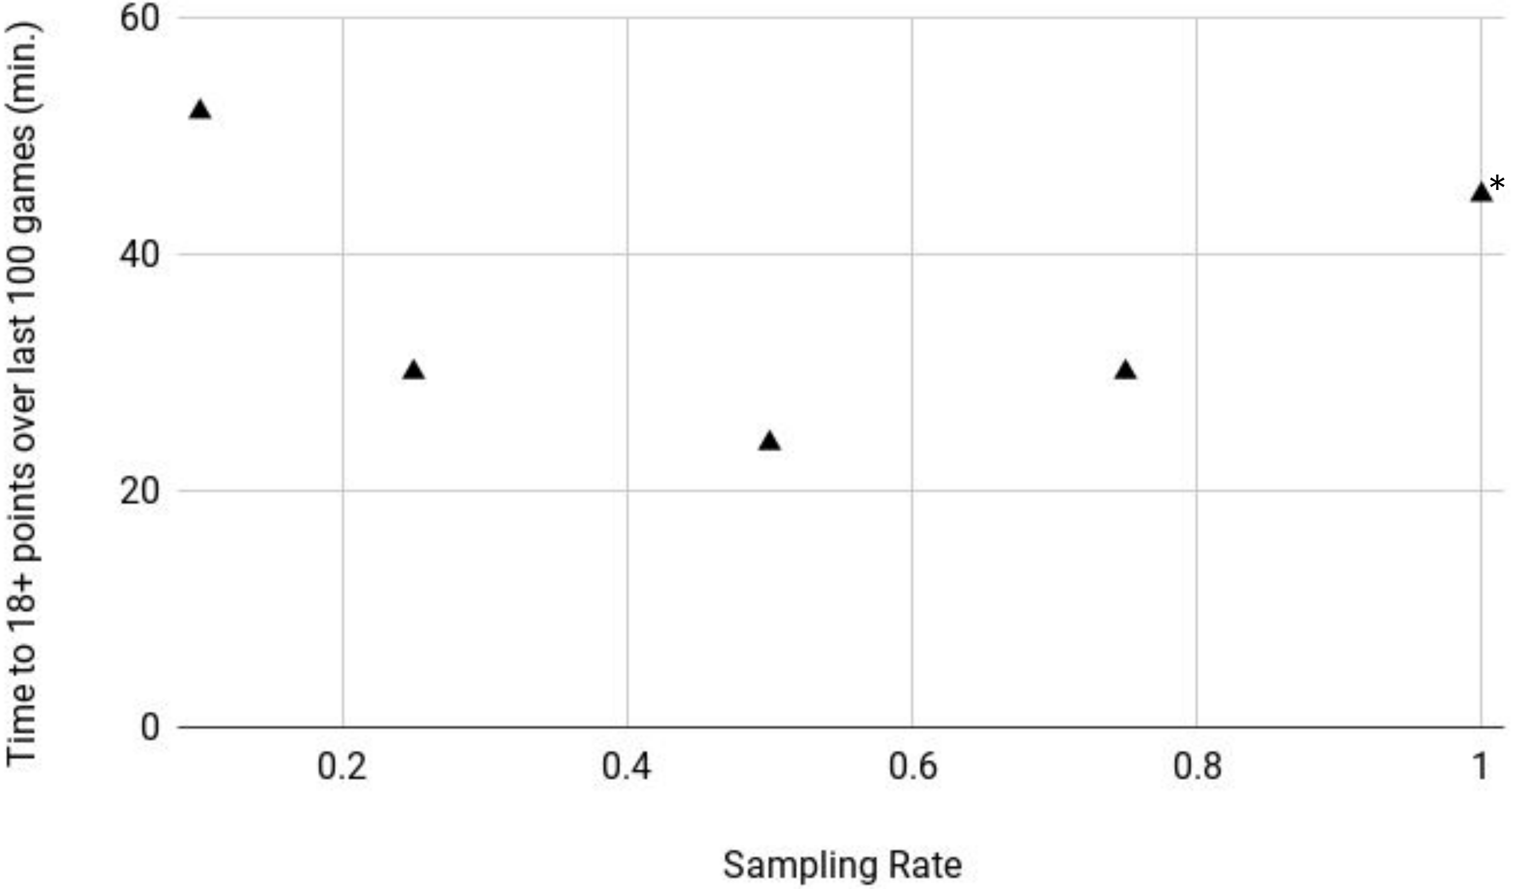
Time to 18+ points over last 100 games (min.)
60
40
20
0
*
0.2
0.4
0.6
0.8
1
Sampling Rate

Results

- Sampling rate of .5 caused best results
- Twice as fast as control network
    - Even when global average pooling (Szegedy et al. 2014) and group normalization (Wu et al. 2018) were added
- Replacing random projection layer with learnable fully-connected layer resulted in slightly worse performance and also took longer to train

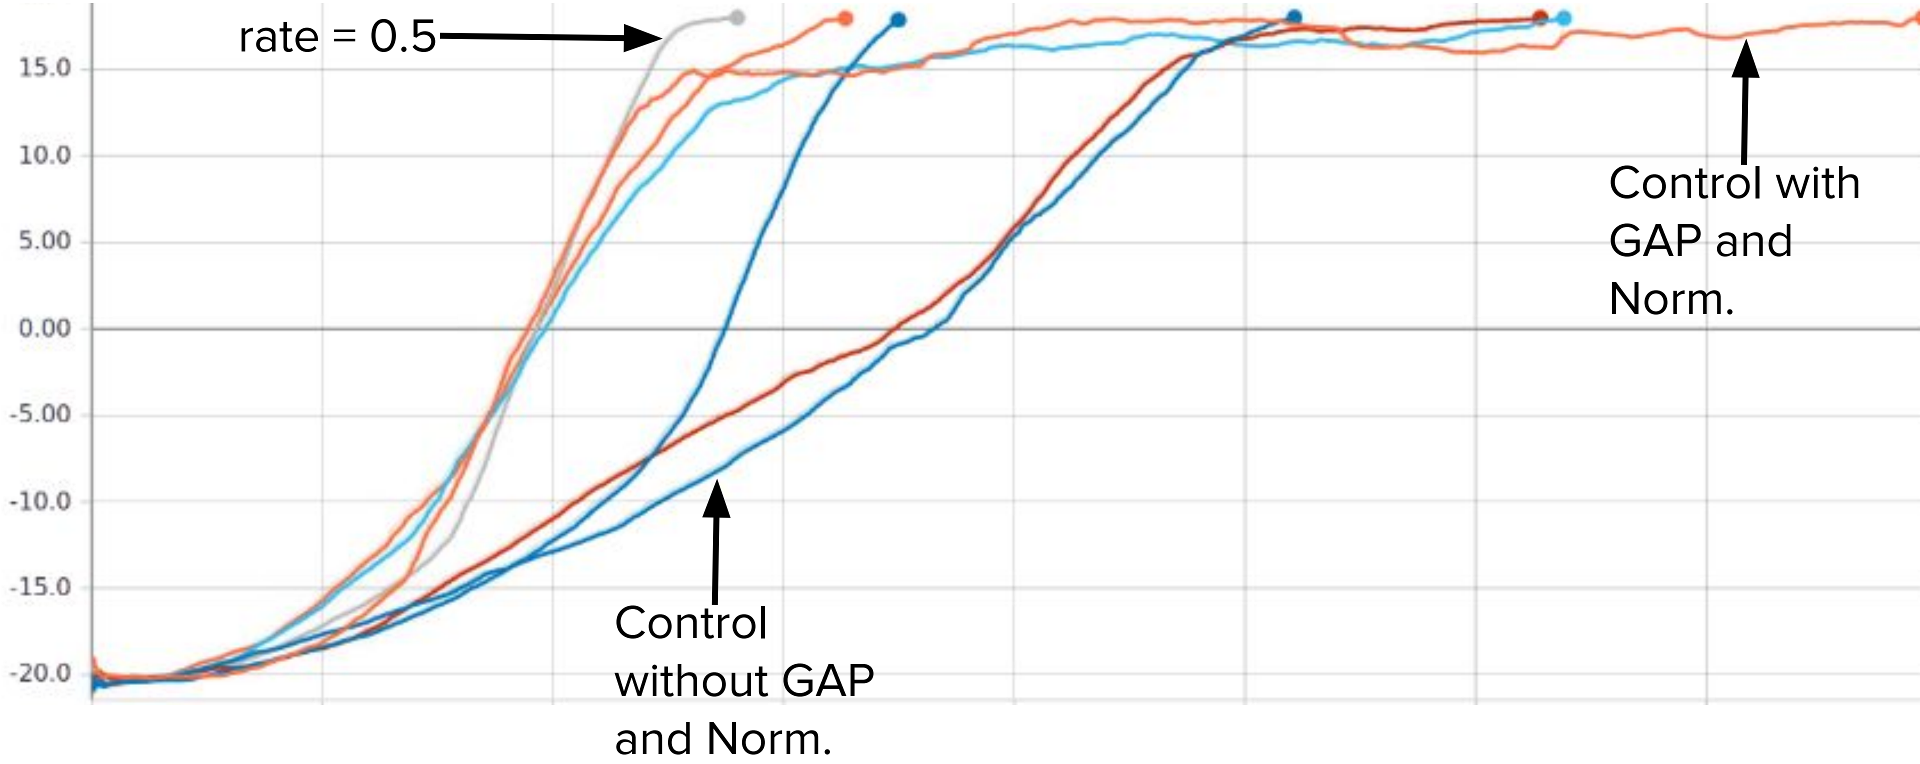
rate = 0.5
15.0
10.0
5.00
0.00
-5.00
-10.0
-15.0
-20.0
Control with GAP and Norm.
Control without GAP and Norm.

Further Investigation

- Random projections applied at *output layer* of AlexNet
  - After global average pooling which was added
- Fully Convolutional
- Trained on MNIST dataset

Classification Results

- The proposed network was competitive with the control
- Examination of activations and weights
  - Much sparser weights and biases - layer 3 shown below

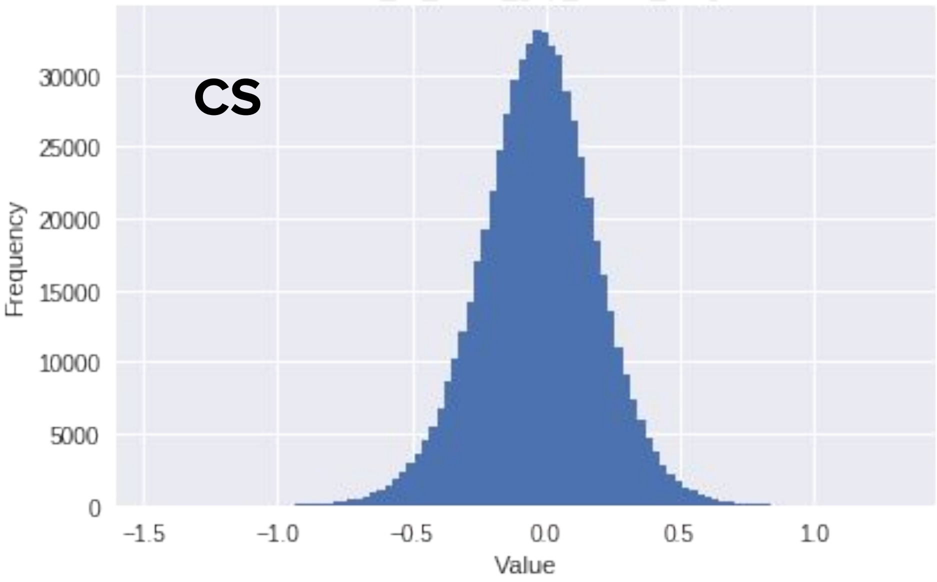

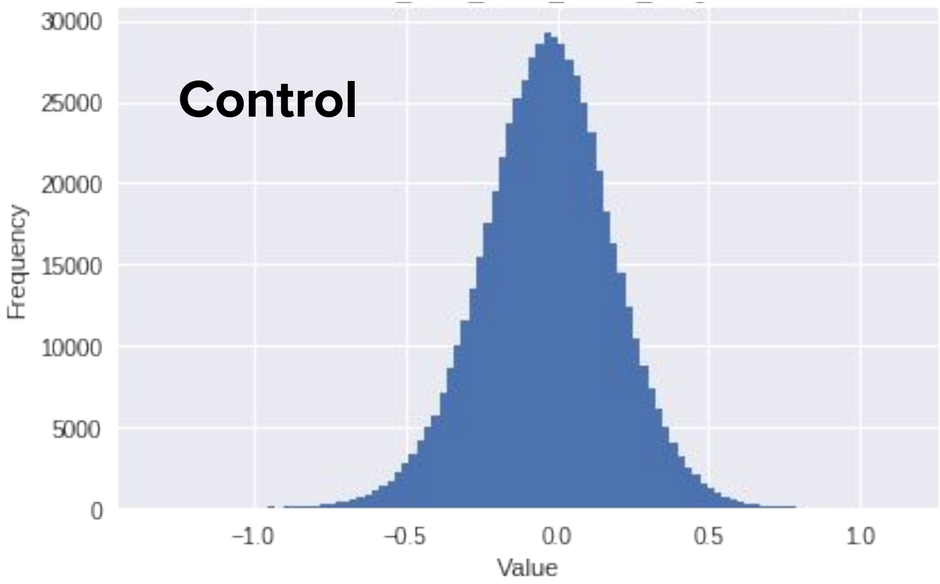

Thank You