LEARNING IMAGE TRANSFORMATIONS WITH CONVOLUTIONAL NEURAL NETWORKS

1) Problem and Objective.

Batch size: # of samples that propagates through the network before updating its weights

- Setting it too high will take too long to converge, and might not fit into memory
- Setting it too low will make the model achieve a local minima.

Are batch sizes affecting the performance of Regression CNNs ?

BY JULIEN LAMON

MARCO LOOG; DAVID TAX; YUKO KATO; TAYLAN TURAN; TOM VIERING; ZIQI WANG

2) Method.

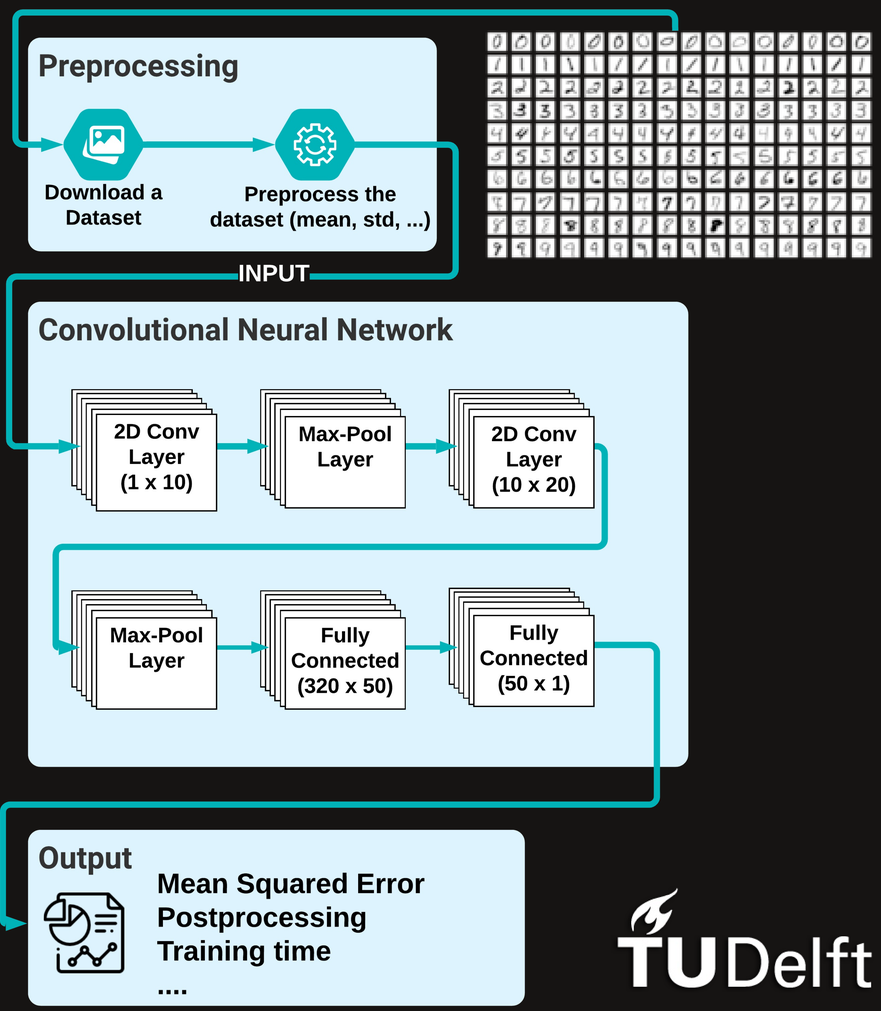

The graphs represents the results of approximating the mean, median and standard deviation of the pixel intensities of the MNIST input images.

- Adam Optimizer is used.
- Learning rate tuned per regression task, on a batch size of 32.
- Model trained on 10 epochs.

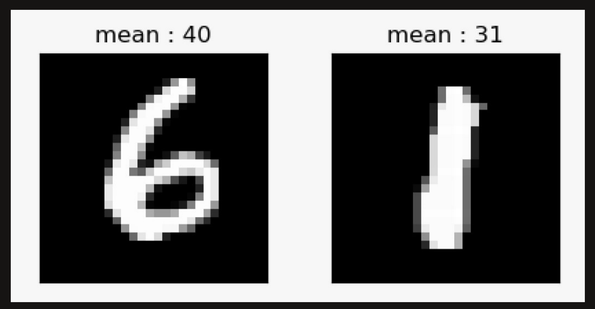

3) Results.

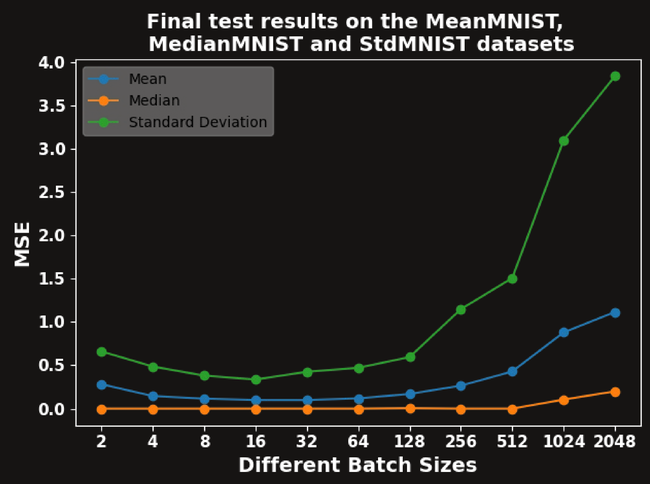

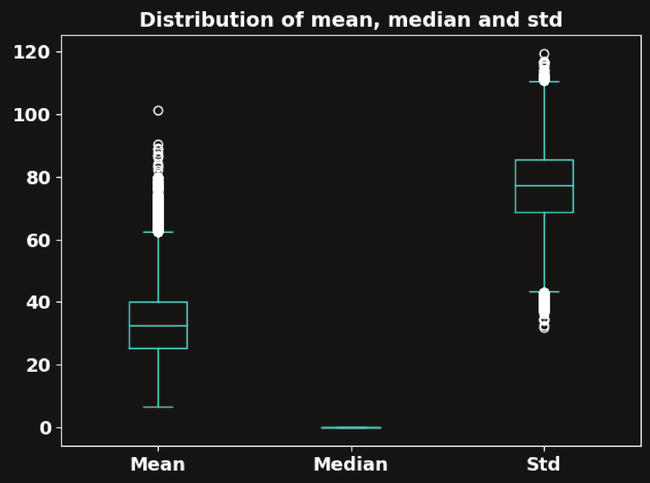

4) Discussion and Conclusion.

- The higher the batch size, the less time it takes
- The higher the batch size, the worse Mean Squared Error it has.
- The optimal batch sizes are found in the interval [8; 32]

Limitations:

- MNIST images are greyscaled. Coloured image will need more parameters, which might affect the performance
- A synthetic dataset created for the research might better indicate how the network is functioning.

TU Delft