




Exploring neural network training dynamics through binary node activations

Gerbrand Haasbroek & Marelle H Davel

Multilingual Speech Technologies, North-West University, South Africa
Centre for Artificial Intelligence Research



- No theoretical framework for DNN training or generalisation
- Generalisation studied using
 - Geometry of loss landscape
 - Stability and robustness
 - Complexity of hypothesis space
 - Margin distributions

- *Sample Set* - Samples that activate a node
- Sample sets formed during the forward pass
- Sample sets refined during the backward pass
- Evolution of sample sets has not been explored

Contributions

- Motivation of the importance of sample sets
- Demonstration of a consistent evolution of sample sets
- Interpretation of findings in terms of clustering

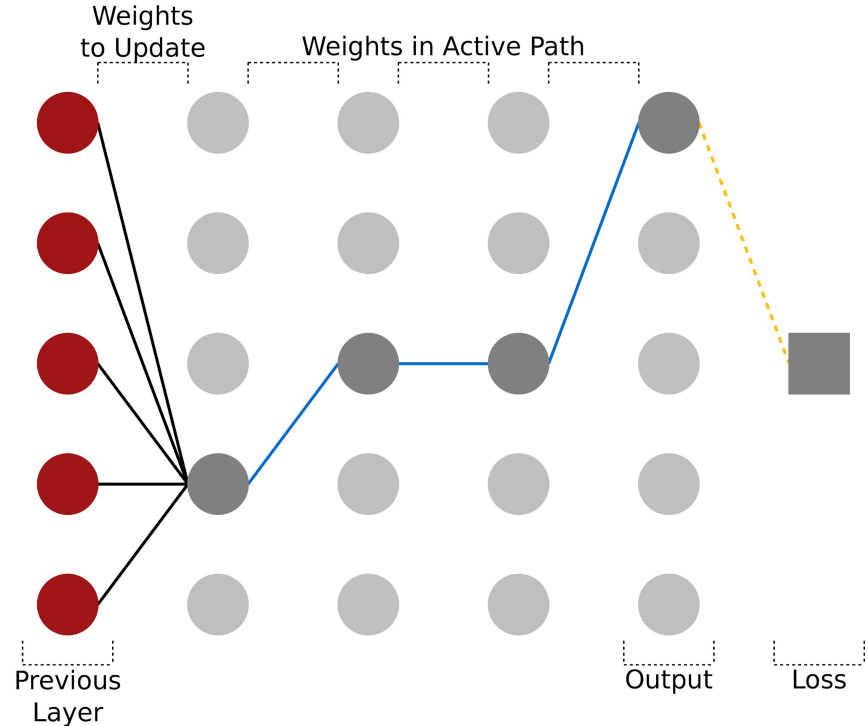
Sample Sets

- Sample added to set if

$$\mathbf{z}_{l-1}^t \cdot \mathbf{w}_{l,j} > \eta \sum_{s \in \hat{S}_{b,l,j}} \left(\mathbf{z}_{l-1}^t \cdot \mathbf{z}_{l-1}^s \right) \phi_{l,j}^s$$

- Sample removed otherwise
- Node-supported cost

$$\phi_{l,j}^s = \sum_{p \in P_j^s} \left(\lambda_p^s \prod_{k=1}^{N-l} w_{p_k} \right)$$




Experiments

- Train fully connected networks:
 - Number of layers: 2, 5, 9
 - Number of nodes per layer: 20, 80, 320
 - Activation functions: ReLU
 - Loss functions: mean squared error, cross entropy
 - Adam optimiser
 - MNIST dataset
 - Epochs: 200

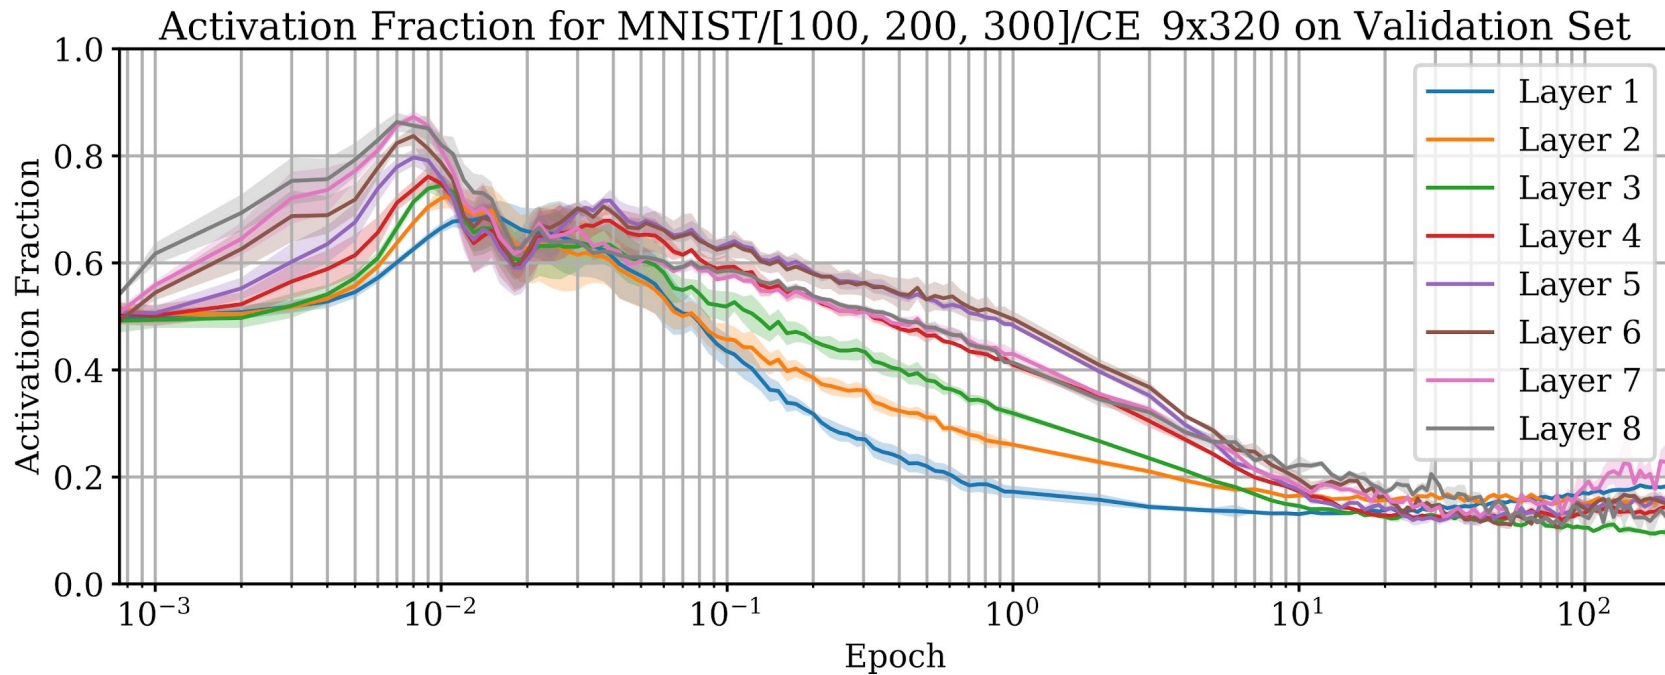
Measurements

- *Activation fraction* - Fraction of all samples in sample set
- *Predictability* - Average amount of information about binary activation given class
- *Informativeness* - Average amount of information about class given binary activation

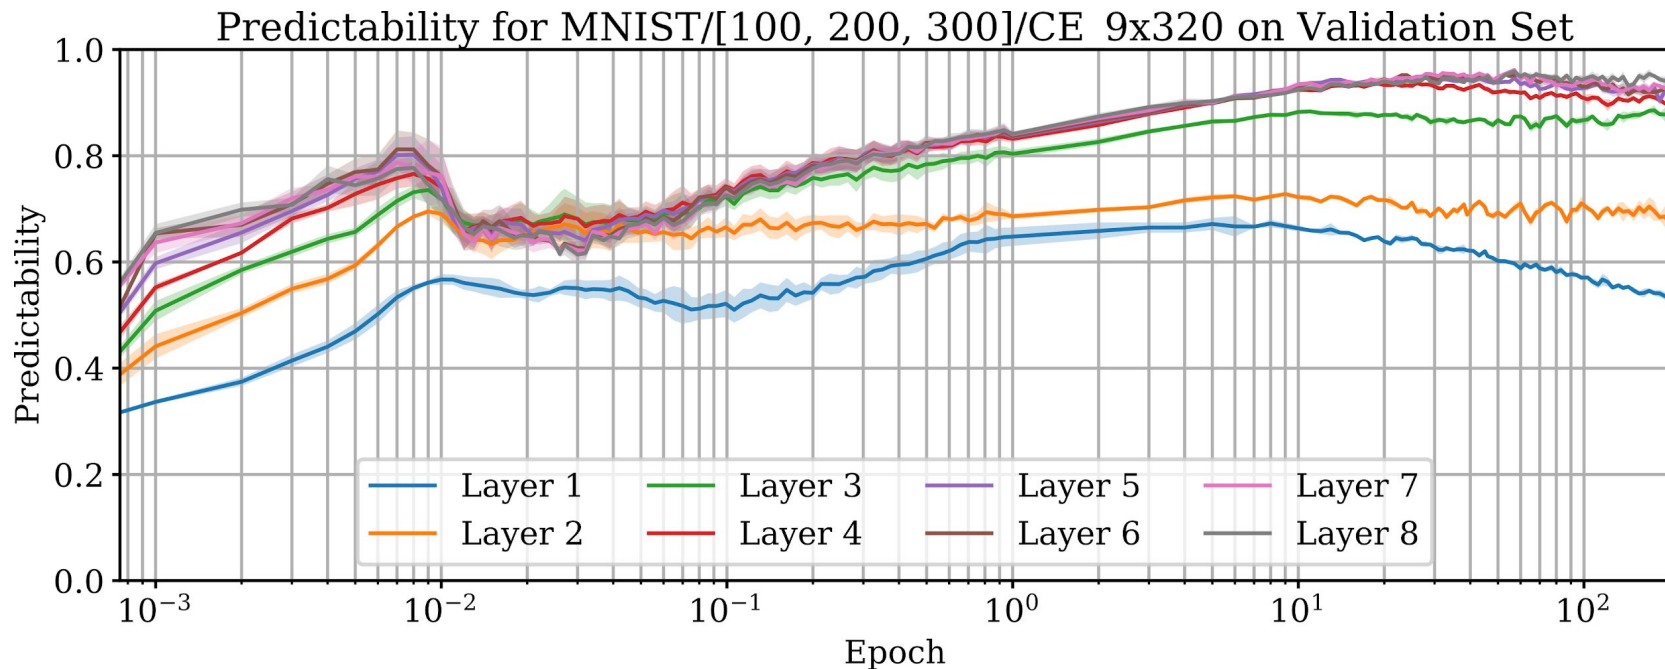
Results



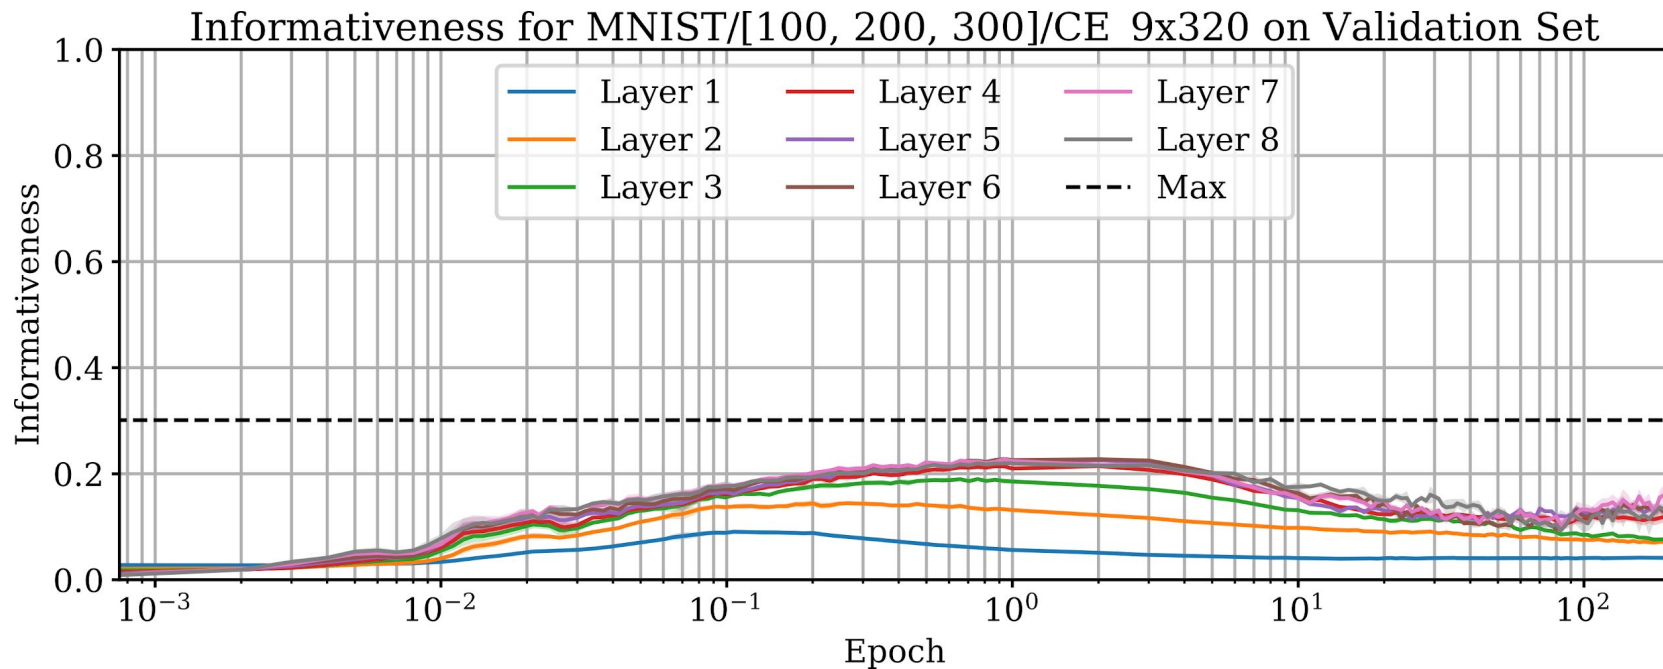
Results



Results




Results



Interpretation

- Initial loss magnitude is large
- Sample sets grow until loss starts decreasing
- Sample sets shrink as loss decreases
- Samples with zero loss do not contribute to weight updates
- Training finds clusters of samples that cause coherent behaviour

Questions



- Predictability

$$p_{l,i} \approx 1 - \frac{1}{\log(2)} H(\hat{Z}_{l,i} | Y)$$

- Informativeness

$$u_{l,i} \approx 1 - \frac{1}{\log(|C|)} H(Y | \hat{Z}_{l,i})$$