Unsupervised Curriculum Learning for Image Clustering

Nathan Hatch Group ID: 10

Discovering the digit ‘7’ is hard without supervision.

Prior Work

Hsu et al. (2016): Deep Image Category Discovery using a Transferred Similarity Function.

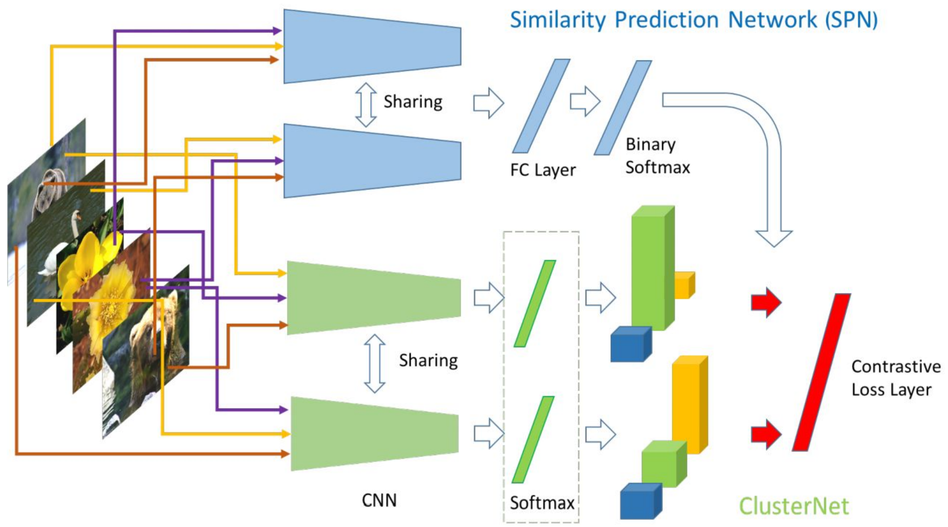

- Similarity Prediction Network: guesses whether input images belong to the same class
  - Transfer learning: trained on Omniglot, applied to MNIST
- Contrastive loss: trains Siamese CNN to cluster images based on similarity

New Idea

Curriculum learning: introduce unlabeled MNIST digits two at a time.

Curriculum	Increment 1	Increment 2	Increment 3	Increment 4
Basic	0-3	0-5	0-7	0-9
Seven First	0-2, 7	0-2, 4-5, 7	0-7	0-9
Seven Last	0-3	0-5	0-6, 8	0-9

Results

Joint training outperforms curriculum learning.

Curriculum	Joint (baseline)	Basic	Seven First	Seven Last
0	99.6	99.6	99.4	99.2
1	98.8	98.8	98.7	98.7
2	99.2	99.1	99.2	98.8
3	98.9	98.5	98.3	98.1
4	96.9	97.0	97.2	97.1
5	99.4	99.1	98.6	98.3
6	99.0	98.8	98.4	98.9
7	95.5	88.7	92.3	87.5
8	99.4	97.1	96.3	96.0
9	98.3	95.4	95.6	96.3
Overall	98.5	97.2	97.4	96.9

The digit ‘7’ is significantly more difficult to discover, especially in curriculum learning.
The ‘continental 7’ is often combined with the digit ‘2’.

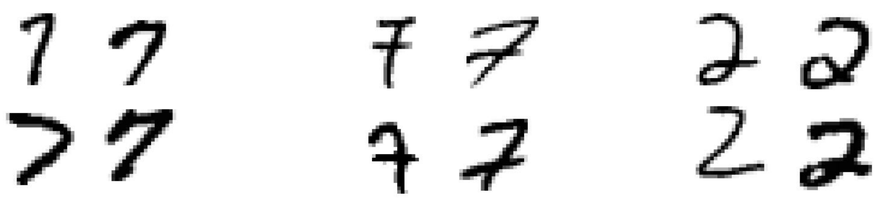

Acknowledgements

Dr. Zsolt Kira, for the concept and for many useful discussions.
Yen-Chang Hsu, for the code base and discussion of prior work.