

DEEP LEARNING WITH KERAS

OCR

Themistoklis Diamantopoulos


Optical Character Recognition

- Handwritten digits
- Classification problem



Feature Representation

- 28 x 28 pixel images
- Pixel intensity between 0 and 1

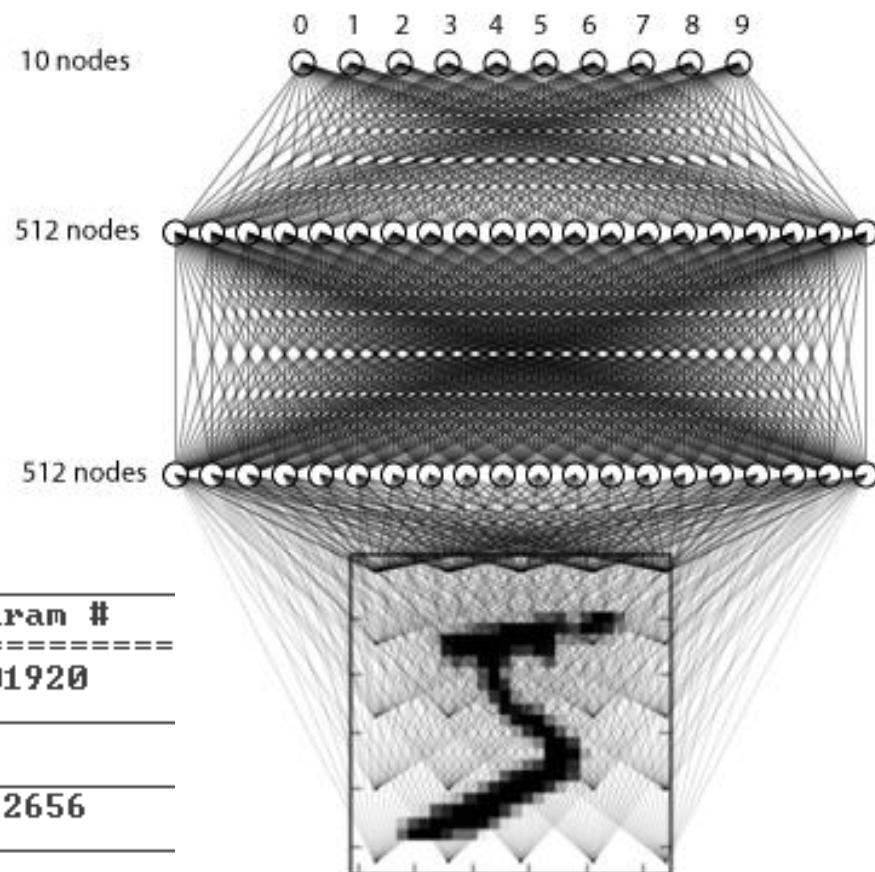


12

0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	.6	.8	0	0	0	0	0	0
0	0	0	0	0	0	.7	1	0	0	0	0	0	0
0	0	0	0	0	0	.7	1	0	0	0	0	0	0
0	0	0	0	0	0	.5	1	.4	0	0	0	0	0
0	0	0	0	0	0	0	1	.4	0	0	0	0	0
0	0	0	0	0	0	0	1	.4	0	0	0	0	0
0	0	0	0	0	0	0	1	.7	0	0	0	0	0
0	0	0	0	0	0	0	1	1	0	0	0	0	0
0	0	0	0	0	0	0	.9	1	.1	0	0	0	0
0	0	0	0	0	0	0	.3	1	.1	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0

Solution using MLP

- 3-layer fully connected network
- Input vector size: 784
- Output layer: 10 nodes
- 2 Intermediate layers




Layer (type)	Output Shape	Param #
dense_1 (Dense)	(None, 512)	401920
dropout_1 (Dropout)	(None, 512)	0
dense_2 (Dense)	(None, 512)	262656
dropout_2 (Dropout)	(None, 512)	0
dense_3 (Dense)	(None, 10)	5130

Solution using CNN


- 2 Conv. layers
- 1 Pooling layer
- Fully Connected Output

Layer (type)	Output Shape	Param #
conv2d_1 (Conv2D)	(None, 26, 26, 32)	320
conv2d_2 (Conv2D)	(None, 24, 24, 32)	9248
max_pooling2d_1 (MaxPooling2)	(None, 12, 12, 32)	0
dropout_1 (Dropout)	(None, 12, 12, 32)	0
flatten_1 (Flatten)	(None, 4608)	0
dense_1 (Dense)	(None, 128)	589952
dropout_2 (Dropout)	(None, 128)	0
dense_2 (Dense)	(None, 10)	1290




Convolution

- Apply kernel to data to provide transformation
- Use padding to preserve input dimensions



Pooling

- Reduce dimensionality, i.e. reduce number of parameters
- Controls overfitting
- Different ways of pooling, e.g. max



Dropout

- Drop some nodes for a pass of the algorithm
- Avoid overtraining certain nodes

