

Neural Networks exam topics

1. Topic
 - a. Local optimization in non-convex cases (reason for non-convexity?)
 - b. RMSprop optimizer
 - c. Dropout
 - d. ResNet
 - e. Gradient ascent
2. Topic
 - a. Weight update strategies
 - b. ReLU and the dying ReLU problem
 - c. LSTM cell
 - d. Convolution as a mathematical operation in continuous and discrete cases
3. Topic
 - a. Newton optimization method
 - b. Ensembling, bagging
 - c. Comparison of loss functions
 - d. Machine learning vs. traditional programming
 - e. Inception
4. Topic
 - a. McCulloch-Pitts model
 - b. Parameters of convolution - filter, stride, padding, etc.
 - c. Linear classifier, margin of the classifier
 - d. Data augmentation
 - e. YOLO
5. Topic
 - a. Statistical learning theory
 - b. Various activation functions and their properties
 - c. Autoencoders
 - d. Graph unrolling and parameter sharing in recurrent neural networks
 - e. MobileNet
6. Topic
 - a. Machine learning problem definition
 - b. Newton optimizer
 - c. Effects and relationship of model capacity and complexity - overfitting, underfitting
 - d. t-Distributed Stochastic Neighbor Embedding
 - e. ShuffleNet
7. Topic
 - a. Credit approval problem
 - b. Objective functions in neural networks
 - c. Nesterov momentum optimizer
 - d. Decomposition of convolutional kernels
 - e. Alexnet + ILSVRC

8. Topic
 - a. Delta learning rule
 - b. Batch normalization
 - c. Transposed convolution, atrous convolution
 - d. Object classification + localization vs. object detection vs. semantic segmentation vs. instance segmentation
 - e. ResNext
9. Topic
 - a. ADAM optimizer
 - b. The softmax function
 - c. R-CNN architectures - R-CNN, Fast R-CNN, Faster R-CNN
 - d. Supervised vs. unsupervised learning
 - e. EfficientNet
10. Topic
 - a. Optimization problem of objective functions of neural networks
 - b. AdaGrad optimizer
 - c. Input vector normalization
 - d. DeconvNet, U-Net
 - e. Neural style transfer
11. Topic
 - a. Multilayer perceptron
 - b. Early stopping
 - c. Gradient descent (multidimensional cases as well)
 - d. Weight regularization (L1, L2)
 - e. Pooling
12. Topic
 - a. Perceptron convergence theorem - no proof required
 - b. Momentum optimizer
 - c. Properties of convolutional neural networks - sparsity, parameter sharing, equivariance, invariance to shifting
 - d. Recurrent neural network examples: predicting the next letter, image captioning
 - e. Adversarial attacks
13. Topic
 - a. Elementary set separation by a single neuron
 - b. Local response normalization
 - c. Unpooling
 - d. Representations - Blum and Li theorem, construction
14. Topic
 - a. Principal component analysis (PCA)
 - b. Back-propagation through time
 - c. Stochastic gradient descent optimizer
 - d. Object detection problem explained
 - e. Effects of filter size on convolution
15. Topic
 - a. Rosenblatt perceptron training algorithm

- b. Back-propagation
- c. Curse of dimensionality
- d. SqueezeNet
- e. Backpropagation and gradient-based optimizers