

IMPORTANT QUESTIONS – DEEP LEARNING (BCS714A)

Module 1: Introducing Deep Learning

1. Explain biological vision and describe the hierarchical organization of visual processing in the human brain with a neat diagram.
2. Explain the experiments of Hubel and Wiesel and describe the role of simple cells and complex cells in visual perception.
3. Explain machine vision and differentiate between traditional machine learning–based vision and deep learning–based vision approaches.
4. Explain the Neocognitron model and describe the functions of S-cells and C-cells with a neat diagram.
5. Explain the LeNet-5 architecture with a neat diagram and discuss its significance in handwritten digit recognition.
6. Explain the ImageNet dataset and the ILSVRC challenge and discuss how they led to the breakthrough of deep learning models such as AlexNet.
7. Explain the AlexNet architecture with a neat diagram and discuss the factors responsible for its breakthrough performance.
8. Explain how deep learning networks learn representations automatically and discuss representation learning with reference to vision and language tasks.
9. Define Natural Language Processing (NLP) and explain the objectives, applications, and brief history of deep learning for NLP.
10. Explain computational representations of language and describe one-hot representations, word vectors, word-vector arithmetic, and localist vs distributed representations with suitable examples.

Module 2: Regularization and Optimization for Deep Learning

1. Explain the need for regularization in deep learning and discuss the problem of overfitting with suitable examples.
2. Explain parameter norm penalties and describe L2 regularization (weight decay) with mathematical formulation.
3. Explain L1 regularization and compare L1 and L2 regularization highlighting their effects on sparsity and feature selection.
4. Explain norm penalties as constrained optimization and describe the role of Lagrange multipliers in regularization.
5. Explain dataset augmentation and discuss how it improves generalization in deep learning models with suitable examples.
6. Explain noise robustness in deep learning and describe dropout as a regularization technique.
7. Explain semi-supervised learning and multi-task learning and discuss how they act as regularizers.

8. Explain early stopping as a regularization technique and describe the algorithm for early stopping with validation error tracking.
9. Explain sparse representations and differentiate between parameter sparsity and representational sparsity.
10. Explain how learning in deep models differs from pure optimization and describe empirical risk minimization.
11. Explain stochastic gradient descent (SGD) and discuss the role of learning rate in optimization.
12. Explain momentum and Nesterov accelerated gradient with intuitive explanation.
13. Explain parameter initialization strategies and discuss Xavier and He initialization methods.
14. Explain algorithms with adaptive learning rates and describe AdaGrad and RMSProp algorithms.

Module 3: Convolutional Neural Networks

1. Explain convolutional neural networks (CNNs) and discuss the motivation behind convolution with respect to sparse interactions, parameter sharing, and equivariant representations.
2. Explain the convolution operation and describe discrete convolution for 1-D and 2-D inputs with a neat diagram.
3. Explain the architecture of a convolutional neural network and describe convolution layer, detector stage, pooling layer, and fully connected layer with a neat diagram.
4. Explain pooling in convolutional neural networks and discuss max pooling and average pooling. Explain how pooling provides translation invariance.
5. Explain convolution and pooling as an infinitely strong prior and discuss the assumptions made by CNNs about image data.
6. Explain variants of the basic convolution function and describe valid, same, full, unshared, and tiled convolution.
7. Explain structured outputs in convolutional networks and discuss pixel-wise labeling and strategies to handle size reduction.
8. Explain different data types handled by convolutional networks and describe convolution on 1-D, 2-D, and 3-D data with examples.
9. Explain efficient convolution algorithms and discuss FFT-based convolution and separable convolution.
10. Explain spatially separable convolution and depthwise separable convolution and compare them with standard convolution.
11. Explain depthwise and pointwise convolution with a neat diagram and discuss their advantages in modern architectures such as MobileNet.
12. Explain the role of convolutional networks in the history of deep learning and discuss their impact on modern vision systems.

Module 4: Recurrent and Recursive Neural Networks (Sequence Modeling)

1. Explain sequence modeling and justify the need for recurrent neural networks over feedforward networks.
2. Explain the structure and working of a basic Recurrent Neural Network (RNN) and describe folded and unfolded representations with a neat diagram.
3. Explain unfolding of computational graphs in recurrent neural networks and describe parameter sharing with suitable example.
4. Explain Backpropagation Through Time (BPTT) and discuss the problems of vanishing and exploding gradients along with gradient clipping as a solution.
5. Explain teacher forcing and discuss the difference between training time and test time behavior of recurrent neural networks.
6. Explain Bidirectional Recurrent Neural Networks and modeling sequences conditioned on context with suitable diagrams and applications.
7. Explain encoder–decoder sequence-to-sequence architecture with a neat diagram and discuss its limitations.
8. Explain the attention mechanism and describe how it overcomes the bottleneck problem in encoder–decoder models.
9. Explain Deep Recurrent Neural Networks and Recursive Neural Networks and differentiate between them with suitable examples.
10. Explain Long Short-Term Memory (LSTM) architecture with a neat diagram and describe how LSTM and GRU networks handle long-term dependencies better than standard RNNs.

Module 5: Interactive Applications of Deep Learning

1. Explain the need for preprocessing natural language data and discuss common preprocessing steps used in NLP.
2. Explain tokenization and discuss sentence tokenization and word tokenization with suitable examples.
3. Explain stop word removal, punctuation removal, and stemming, and discuss their impact on NLP model performance.
4. Explain n-grams and describe the handling of bigrams and trigrams in natural language preprocessing.
5. Explain preprocessing of the full corpus and discuss how preprocessing decisions vary with dataset size.
6. Explain word embeddings and discuss their advantages over one-hot and bag-of-words representations.
7. Explain the theory behind word2vec and describe the distributional hypothesis.
8. Explain the CBOW and Skip-gram architectures of word2vec and compare them.
9. Explain intrinsic and extrinsic evaluation methods for word vectors and discuss their trade-offs.
10. Explain visualization of word vectors using dimensionality reduction techniques such as t-SNE and discuss its usefulness.
11. Explain the confusion matrix, ROC curve, and AUC, and discuss their significance in evaluating NLP models.

12. Explain sentiment classification using a dense neural network and describe the role of embedding and dense layers.
13. Explain the convolutional neural network approach for text classification and discuss how convolution captures local word patterns.
14. Explain the use of recurrent neural networks, LSTM, and Bidirectional LSTM for natural language classification.
15. Explain sequence-to-sequence models with attention and discuss their applications along with the role of transfer learning in NLP.