SIDDARTH INSTITUTE OF ENGINEERING & TECHNOLOGY::PUTTUR
(AUTONOMOUS)
(Approved by AICTE, New Delhi & Affiliated to JNTUA, Ananthapuramu)
Puttur -517583, Tirupathi District, A.P. (India)

QUESTION BANK (DESCRIPTIVE)

Subject with Code: DEEP LEARNING (23CS0528)
Year & Sem: IV B.Tech & I Sem

Course & Branch: CSE
Regulation: R23

UNIT - I: Linear Algebra, Probability & Information Theory, Numerical Computation

1	a)	Define scalar, vector, matrix and tensor with suitable examples.	[L1,CO1]	[2M]
	b)	List the different types of matrices used in linear algebra.	[L1,CO1]	[2M]
	c)	Define eigenvalues and eigenvectors of a matrix.	[L1,CO1]	[2M]
	d)	State Bayes' Rule.	[L1,CO1]	[2M]
	e)	Define expectation, variance and covariance of a random variable.	[L1,CO1]	[2M]
2	a)	Explain the basic matrix operations (addition, multiplication, transpose, inverse) with examples.	[L2,CO1]	[5M]
	b)	Explain the concept of norms and their significance in machine learning.	[L2,CO1]	[5M]
3		Explain Eigen Decomposition and Singular Value Decomposition (SVD) with an example.	[L5,CO1]	[10M]
4	a)	Discuss Principal Component Analysis (PCA) and its role in dimensionality reduction.	[L4,CO1]	[5M]
	b)	Differentiate between marginal probability and conditional probability with examples.	[L4,CO1]	[5M]
5	a)	Illustrate, with a numerical example, the problems of overflow and underflow in numerical computation.	[L4,CO1]	[5M]
	b)	Analyze the working of gradient-based optimization algorithms.	[L4,CO1]	[5M]
6	a)	Explain constrained optimization and the method of Lagrange multipliers with an example.	[L2,CO1]	[5M]
	b)	Analyze the Linear Least Squares method for solving a regression problem.	[L4,CO1]	[5M]
7		Explain the fundamentals of Information Theory (self-information, entropy, KL divergence, cross-entropy) with their mathematical formulations.	[L5,CO1]	[10M]
8		Apply Singular Value Decomposition to reduce the dimensionality of a given dataset. Illustrate with an example.	[L3,CO1]	[10M]
9		Evaluate the significance of common probability distributions in machine learning, justifying your answer with examples.	[L6,CO1]	[10M]
10		Solve a numerical problem to compute the eigenvalues and eigenvectors of a given matrix, and explain how this relates to PCA computation.	[L5,CO5]	[10M]
11	a)	Define random variable and explain the types of probability distributions.	[L1,CO1,CO5]	[5M]
	b)	Explain gradient-based optimization techniques used to minimize a cost function.	[L4,CO1]	[5M]

UNIT - II: Machine Learning Basics & Deep Feedforward Networks

1	a)	Define overfitting and underfitting.	[L1,CO2]	[2M]
	b)	What are hyperparameters? Give examples.	[L1,CO2]	[2M]
	c)	Define Maximum Likelihood Estimation.	[L1,CO2]	[2M]
	d)	List commonly used hidden unit activation functions.	[L1,CO2]	[2M]
	e)	Define bias and variance of an estimator.	[L1,CO2]	[2M]
2	a)	Explain the difference between supervised and unsupervised learning with examples.	[L2,CO2]	[5M]
	b)	Explain the role of validation sets in hyperparameter selection and model evaluation.	[L2,CO2]	[5M]
3		Explain the Bias-Variance tradeoff and its impact on model performance.	[L5,CO2]	[10M]
4	a)	Discuss the challenges that motivate the need for deep learning over traditional machine learning.	[L4,CO2]	[5M]
	b)	Illustrate how a feedforward network learns the XOR function.	[L4,CO2]	[5M]
5	a)	Design a simple feedforward neural network architecture for a given classification problem.	[L4,CO2]	[5M]
	b)	Analyze the working of the Stochastic Gradient Descent algorithm.	[L4,CO2]	[5M]
6	a)	Explain Bayesian statistics and how it differs from the frequentist approach.	[L2,CO2]	[5M]
	b)	Analyze the Back-Propagation algorithm with a suitable example.	[L4,CO2]	[5M]
7		Justify Gradient-Based Learning in neural networks, including cost function formulation and output units.	[L5,CO2]	[10M]
8		Apply the concept of estimators to explain how model parameters are learned from data.	[L3,CO2]	[10M]
9		Evaluate different architecture design choices (depth, width, number of hidden units) and their effect on network performance.	[L6,CO2]	[10M]
10		Solve the XOR problem using a feedforward network and derive the back-propagation equations for updating its weights.	[L5,CO2]	[10M]
11	a)	Define estimator, bias and consistency of an estimator.	[L1,CO2]	[5M]
	b)	Explain the role of hidden units and activation functions in deep feedforward networks.	[L4,CO2]	[5M]

UNIT - III: Regularization for Deep Learning & Optimization for Training Deep Models

1	a)	Define regularization.	[L1,CO3]	[2M]
	b)	What is Dropout?	[L1,CO3]	[2M]
	c)	Define Early Stopping.	[L1,CO3]	[2M]
	d)	List the types of parameter norm penalties.	[L1,CO3]	[2M]
	e)	Define Bagging and Ensemble methods.	[L1,CO3]	[2M]
2	a)	Explain L1 and L2 parameter norm penalties with their mathematical formulations.	[L2,CO3]	[5M]
	b)	Explain Dataset Augmentation and Noise Robustness as regularization techniques.	[L2,CO3]	[5M]
3		Explain the Dropout technique and how it helps prevent overfitting in deep networks.	[L5,CO3]	[10M]
4	a)	Discuss Semi-Supervised Learning and Multi-Task Learning approaches.	[L4,CO3]	[5M]
	b)	Illustrate Parameter Tying and Parameter Sharing with an example (e.g., CNN weight sharing).	[L4,CO3]	[5M]
5	a)	Design a training strategy that uses Early Stopping to prevent overfitting.	[L4,CO3]	[5M]
	b)	Analyze the challenges in neural network optimization (local minima, saddle points, ill-conditioning).	[L4,CO3]	[5M]
6	a)	Explain parameter initialization strategies and their importance in training deep models.	[L2,CO3]	[5M]
	b)	Analyze adaptive learning rate algorithms such as AdaGrad, RMSProp and Adam.	[L4,CO3]	[5M]
7		Explain Adversarial Training and Tangent Prop / Manifold Tangent Classifier methods.	[L5,CO3]	[10M]
8		Apply approximate second-order optimization methods (Newton's method, Conjugate Gradient) to a given optimization problem.	[L3,CO3]	[10M]
9		Evaluate the effectiveness of various optimization strategies and meta-algorithms (batch normalization, coordinate descent) used in training deep models.	[L6,CO3]	[10M]
10		Solve a numerical example demonstrating how Bagging and Ensemble methods reduce the variance of a model's predictions.	[L5,CO3]	[10M]
11	a)	Define sparse representation and explain its advantages.	[L1,CO3]	[5M]
	b)	Explain the challenges involved in pure optimization for training deep models.	[L4,CO3]	[5M]

UNIT - IV: Convolutional Networks

1	a)	Define the convolution operation.	[L1,CO4]	[2M]
	b)	What is pooling? List its types.	[L1,CO4]	[2M]
	c)	Define structured output in the context of CNNs.	[L1,CO4]	[2M]
	d)	List the different data types that can be processed using convolution.	[L1,CO4]	[2M]
	e)	Define feature map and receptive field.	[L1,CO4]	[2M]
2	a)	Explain the convolution operation with a mathematical example.	[L2,CO4]	[5M]
	b)	Explain basic convolution functions with respect to stride, padding and dilation.	[L2,CO4]	[5M]
3		Explain the architecture and working of a Convolutional Neural Network in detail.	[L5,CO4]	[10M]
4	a)	Discuss the different types of pooling operations (max, average, global).	[L4,CO4]	[5M]
	b)	Illustrate how structured outputs are generated for tasks such as image segmentation.	[L4, CO6]	[5M]
5	a)	Design a CNN architecture for a given image classification task.	[L4,CO6]	[5M]
	b)	Analyze efficient convolution algorithms (FFT-based, Winograd) and their advantages.	[L4,CO4]	[5M]
6	a)	Explain how CNNs handle different data types (1-D, 2-D and 3-D data).	[L2,CO4]	[5M]
	b)	Analyze the use of random or unsupervised features in convolutional networks.	[L4,CO4]	[5M]
7		Explain the theoretical basis for convolutional networks: sparse interactions, parameter sharing and equivariant representations.	[L5,CO4]	[10M]
8		Apply convolution and pooling operations manually on a sample input matrix with a given filter.	[L3,CO4]	[10M]
9		Evaluate the performance trade-offs between different convolution algorithms in terms of computational efficiency.	[L6,CO4]	[10M]
10		Solve a numerical example to compute the output feature map dimensions given input size, filter size, stride and padding.	[L5,CO4]	[10M]
11	a)	Define equivariance to translation in the context of CNNs.	[L1,CO4]	[5M]
	b)	Explain the role of pooling layers in achieving approximate translation invariance.	[L4,CO4]	[5M]

UNIT - V: Sequence Modelling - Recurrent and Recursive Nets

1	a)	Define Recurrent Neural Network (RNN).	[L1,CO5]	[2M]
	b)	What is an autoencoder?	[L1,CO5]	[2M]
	c)	Define LSTM and explain the need for it.	[L1,CO5]	[2M]
	d)	List the types of gated RNN units.	[L1,CO5]	[2M]
	e)	Define the Encoder-Decoder architecture.	[L1,CO5]	[2M]
2	a)	Explain the concept of unfolding computational graphs in RNNs.	[L2,CO5]	[5M]
	b)	Explain Bidirectional RNNs with a suitable diagram.	[L2,CO5]	[5M]
3		Explain the architecture and working of LSTM networks in detail.	[L5,CO5]	[10M]
4	a)	Discuss Deep Recurrent Networks and Recursive Neural Networks.	[L4,CO5]	[5M]
	b)	Illustrate the Encoder-Decoder sequence-to-sequence architecture with an example such as machine translation.	[L4,CO5]	[5M]
5	a)	Design a sequence model for a given time-series prediction problem using LSTM/GRU.	[L4,CO5]	[5M]
	b)	Analyze the problem of long-term dependencies in RNNs and techniques used to address it.	[L4,CO5]	[5M]
6	a)	Explain Echo State Networks and their working principle.	[L2,CO5]	[5M]
	b)	Analyze the differences between GRU and LSTM gated units.	[L4,CO5]	[5M]
7		Explain deep generative models (e.g., autoencoders, VAEs, GANs) and their applications in sequence modelling.	[L5,CO5]	[10M]
8		Apply an autoencoder to a given dimensionality reduction / reconstruction problem.	[L3,CO5]	[10M]
9		Evaluate the effectiveness of LSTM over vanilla RNN in handling the vanishing gradient problem, justifying your answer.	[L6,CO5]	[10M]
10		Derive the forward and backward pass equations for a simple RNN cell.	[L5,CO5]	[10M]
11	a)	Define the vanishing and exploding gradient problems in RNNs.	[L1,CO5]	[5M]
	b)	Explain optimization strategies used for handling long-term dependencies in sequence models.	[L4,CO5]	[5M]