



SNS COLLEGE OF ENGINEERING
Kurumbapalayam (Po), Coimbatore – 641 107
AN AUTONOMOUS INSTITUTION



Approved by AICTE, New Delhi & Affiliated to Anna University, Chennai
DEPARTMENT OF ARTIFICIAL INTELLIGENCE AND DATA SCIENCE
UNIT -4

2 marks:

1. What is the difference between convex and non-convex optimization?

- Convex optimization has a single global minimum, while non-convex optimization has multiple local minima and saddle points.

2. Why is optimization in deep learning often non-convex?

- Deep neural networks have complex, high-dimensional loss surfaces with multiple local minima and saddle points, making them inherently non-convex.

3. What are saddle points, and how do they affect optimization in deep learning?

- Saddle points are points where the gradient is zero but are neither a minimum nor maximum. They slow down optimization as gradient-based methods struggle to escape them.

4. How does gradient descent help in non-convex optimization?

- Gradient descent iteratively updates weights in the direction of the negative gradient, helping find an optimal or near-optimal solution even in non-convex loss landscapes.

5. What is stochastic gradient descent (SGD)?

- SGD updates model parameters using a small random subset (mini-batch) of data instead of the full dataset, making optimization faster and more efficient.

6. How does mini-batch gradient descent differ from full-batch gradient descent?

- Mini-batch gradient descent updates weights using small subsets of data, while full-batch uses the entire dataset for each update. Mini-batch is faster and generalizes better.

7. What is the role of the learning rate in SGD?

- The learning rate controls the step size of weight updates. A small learning rate leads to slow convergence, while a large one can cause instability.

8. What is momentum in the context of SGD?

- Momentum helps accelerate SGD by smoothing updates using a moving average of past gradients, reducing oscillations and improving convergence speed.

9. What is overfitting, and how can it be reduced in deep learning?

- Overfitting occurs when a model learns patterns specific to training data but fails to generalize to new data. It can be reduced using regularization, dropout, and data augmentation.

10. How does dropout improve generalization in neural networks?

- Dropout randomly deactivates neurons during training, forcing the network to learn more robust features and reducing overfitting.
11. What is the role of weight regularization in improving generalization?
 - Weight regularization techniques like L1 (Lasso) and L2 (Ridge) regularization prevent large weight values, reducing overfitting and improving generalization.
 12. How does data augmentation help in generalization?
 - Data augmentation artificially increases dataset size by applying transformations (rotation, flipping, cropping) to input images, improving model robustness.
 13. What is a Spatial Transformer Network (STN), and why is it useful?
 - STN is a neural network module that learns spatial transformations, allowing CNNs to be invariant to scale, rotation, and distortion.
 14. How does an STN differ from a conventional CNN?
 - Unlike CNNs, STNs actively transform input data to align features, improving performance on tasks requiring spatial invariance.
 15. What are the main components of an STN?
 - Localization network, grid generator, and sampler. These components work together to transform the input dynamically.
 16. How does the localization network in STNs work?
 - It predicts transformation parameters from input features and applies an affine transformation to align the input properly.
 17. What is the vanishing gradient problem in RNNs?
 - When training RNNs with backpropagation through time (BPTT), gradients shrink exponentially, making it hard to learn long-term dependencies.
 18. How does an LSTM mitigate the vanishing gradient problem?
 - LSTMs use gates (input, forget, and output) to control information flow, allowing them to retain long-term dependencies.
 19. What is the main difference between an LSTM and a GRU?
 - GRUs have fewer gates (update and reset) than LSTMs, making them computationally faster while still handling long-term dependencies well.
 20. How does a Word-Level RNN differ from a Character-Level RNN?
 - Word-level RNNs process words as input tokens, while character-level RNNs process individual characters, making them suitable for text generation and spell-checking.

16 marks:

1. Describe the concept of optimization in deep learning and outline its various types.

Optimization in Deep Learning:

In Deep Learning, with the help of loss function, the performance of the model is estimated/ evaluated. This loss is used to train the network so that it performs better. Essentially, we try to minimize the Loss function. Lower Loss means the model performs

better. The Process of minimizing any mathematical function is called Optimization.

Optimizers are algorithms or methods used to change the features of the neural network such as weights and learning rate so that the loss is reduced. Optimizers are used to solve optimization problems by minimizing the function

The Goal of an Optimizer is to minimize the Objective Function(Loss Function based on the Training Data set). Simply Optimization is to minimize the Training Error.

Need for Optimization:

- Prescence of Local Minima reduces the model performance
- Prescence of Saddle Points which creates Vanishing Gradients or Exploding Gradient Issues
- To select appropriate weight values and other associated model parameters
- To minimize the loss value (Training error)

Convex Optimization:

Convex optimization is a kind of optimization which deals with the study of problem of minimizing convex functions. Here the optimization function is convex function.

All Linear functions are convex, so linear programming problems are convex problems. When we have a convex objective and a convex feasible region, then there can be only one optimal solution, which is globally optimal.

Definition: A set $C \subseteq \mathbb{R}^n$ is convex if for $x, y \in C$ and any $\alpha \in [0, 1]$,

$$\alpha x + (1 - \alpha)y \in C$$

Convexity plays a vital role in the design of optimization algorithms. This is largely due to the fact that it is much easier to analyze and test algorithms in such a context.

Consider the given Figure 4.1 given below, select any to points in the region and join them by a straight Line. If the line and the selected points all lie inside the region then we call that region as Convex Region (as Shown in the diagram Figure 4.1)

Non-Convex Optimization:

- The Objective function is a non- convex function
- All non-linear problems can be modelled by using non convex functions. (Linear functions are convex)
- It has Multiple feasible regions and multiple locally optimal points.
- There can't be a general algorithm to solve it efficiently in all cases
- Neural networks are universal function approximators, to do this, they need to be able to approximate non-convex functions.

Refer the figure 4.2 .It shows Non Convex Region

How to solve non-convex problems?

- ★ Stochastic gradient descent
- ★ Mini-batching
- ★ SVRG
- ★ Momentum

Reasons For Non-Convexity:

- Presence of many Local Minima
- Presence of Saddle Points
- Very Flat Regions
- Varying Curvature

2. What is a Spatial Transformer Network (STN), and what is its significance?

Spatial Transform Network [STN]:

Spatial Transformer Network (STN) helps to crop out and scale-normalizes the appropriate region, which can simplify the subsequent classification task and lead to better classification performance. The Spatial Transformer Network contains three parts Namely, Localization, Grid Generator and Sampler. These Networks are used for performing Transformations such as Cropping, Rotation etc on the given input images.

Localisation Net:

With input feature map U , with width W , height H and C channels, outputs are θ , the parameters of transformation $T\theta$. It can be learnt as affine transform

Grid Generator:

Suppose we have a regular grid G , this G is a set of points with target coordinates (x_{t_i}, y_{t_i}) . Then we apply transformation $T\theta$ on G , i.e. $T\theta(G)$. After $T\theta(G)$, a set of points with destination coordinates (x_{t_i}, y_{t_i}) is outputted. These points have been altered based on the transformation parameters. It can be Translation, Scale, Rotation or More Generic Warping depending on how we set θ as mentioned above.

Sampler:

Based on the new set of coordinates (x_{t_i}, y_{t_i}) , we generate a transformed output feature map V . This V is translated, scaled, rotated, warped, projective transformed or affined, whatever. It is noted that STN can be applied to not only input image, but also intermediate feature maps.

➤ STN is a mechanism that rotates or scales an input image or a feature map in order to focus on the target object and to remove rotational variance .

❑ One of the most notable features of STNs is their modularity (the module can be injected into any part of the model) and their ability to be trained with a single backprop algorithm without modification of the initial model.

Advantages:

- ❖ Helps in learning explicit spatial transformations like translation, rotation, scaling, cropping, non-rigid deformations, etc. of features.
- ❖ Can be used in any networks and at any layer and learnt in an end-to-end trainable manner.
- ❖ Provides improvement in the performance of existing models.

3.Explain the structure and working mechanism of Long Short-Term Memory (LSTM) networks.

Long Short Term Memory Network's (LSTM):

LSTMs are a special kind of RNN — capable of learning long-term dependencies by remembering information for long periods is the default behavior. All RNN are in the form of a chain of repeating modules of a neural network. In standard RNNs, this repeating module will have a very simple structure, such as a single tanh layer.

LSTMs also have a chain-like structure, but the repeating module is a bit different structure. Instead of having a single neural network layer, four interacting layers are communicating extraordinarily.

Hochreiter & Schmidhuber (1997) solved the problem of getting an RNN to remember things for a long time (like hundreds of time steps). They designed a memory cell using logistic and linear units with multiplicative interactions. Information gets into the cell whenever its “write” gate is on. The information stays in the cell so long as its “keep” gate is on. Information can be read from the cell by turning on its “read” gate.(Refer Figure 4.6 – shown Below)

To preserve information for a long time in the activities of an RNN, we use a circuit that implements an analog memory cell.

- A linear unit that has a self-link with a weight of 1 will maintain its state. – Information is stored in the cell by activating its write gate.
- Information is retrieved by activating the read gate.
- We can backpropagate through this circuit because logistics are had nice derivatives.

Steps Involved in LSTM Networks:

Step 1: Decide how much past data it should remember

The first step in the LSTM is to decide which information should be omitted from the cell in that particular time step. The sigmoid function determines this. It looks at the previous state (h_{t-1}) along with the current input x_t and computes the function.

Step 2: Decide how much this unit adds to the current state

In the second layer, there are two parts. One is the sigmoid function, and the other is the tanh function. In the sigmoid function, it decides which values to let through (0 or 1). tanh function gives weightage to the values which are passed, deciding their level of importance (-1 to 1).

Step 3: Decide what part of the current cell state makes it to the output

The third step is to decide what the output will be. First, we run a sigmoid layer, which decides what parts of the cell state make it to the output. Then, we put the cell state through tanh to push the values to be between -1 and 1 and multiply it by the output of the sigmoid gate.

Applications of LSTM include:

- Robot control
- Time series prediction
- Speech recognition
- Rhythm learning
- Music composition
- Grammar learning
- Handwriting recognition

4. Define Computational Neuroscience and Artificial Neuroscience, highlighting their connection to deep learning.

Computational and Artificial Neuro-Science:

Computational neuroscience is the field of study in which mathematical tools and theories are used to investigate brain function.

The term “computational neuroscience” has two different definitions:

1. using a computer to study the brain
2. studying the brain as a computer

Computational and Artificial Neuroscience deals with the study or understanding of

how signals are transmitted through and from the human brain. A better understanding of How decision is made in human brain by processing the data or signals will help us in developing Intelligent algorithms or programs to solve complex problems. Hence, we need to understand the basics of Biological Neural Networks (BNN).

The Biological Neurons:

The human brain consists of a large number, more than a billion of neural cells that process information. Each cell works like a simple processor. The massive interaction between all cells and their parallel processing only makes the brain's abilities possible. Figure 1 represents a human biological nervous unit. Various parts of biological neural network(BNN) is marked in Figure 4.7.

Dendrites are branching fibres that extend from the cell body or soma.

Soma or cell body of a neuron contains the nucleus and other structures, support chemical processing and production of neurotransmitters.

Axon is a singular fiber carries information away from the soma to the synaptic sites of other neurons (dendrites and somas), muscles, or glands.

Axon hillock is the site of summation for incoming information. At any moment, the collective influence of all neurons that conduct impulses to a given neuron will determine whether or not an action potential will be initiated at the axon hillock and propagated along the axon.

Myelin sheath consists of fat-containing cells that insulate the axon from electrical activity. This insulation acts to increase the rate of transmission of signals. A gap exists between each myelin sheath cell along the axon. Since fat inhibits the propagation of electricity, the signals jump from one gap to the next.

Nodes of Ranvier are the gaps (about 1 μm) between myelin sheath cells. Since fat serves as a good insulator, the myelin sheaths speed the rate of transmission of an electrical impulse along the axon.

Synapse is the point of connection between two neurons or a neuron and a muscle or a gland. Electrochemical communication between neurons take place at these junctions.

Terminal buttons of a neuron are the small knobs at the end of an axon that release chemicals called neurotransmitters.

Information flow in a neural cell

The input/output and the propagation of information are shown below.

Artificial neuron model

An artificial neuron is a mathematical function conceived as a simple model of a real (biological) neuron.

- The McCulloch-Pitts Neuron

This is a simplified model of real neurons, known as a Threshold Logic Unit. • A set of input connections brings in activations from other neurons.

- A processing unit sums the inputs, and then applies a non-linear activation function (i.e. squashing/transfer/threshold function).
- An output line transmits the result to other neurons.

Basic Elements of ANN:

Neuron consists of three basic components –weights, thresholds and a single activation function. An Artificial neural network(ANN) model based on the biological neural systems is shown in figure 4.8.

The goal of computational neuroscience is to explain how electrical and chemical signals are used in the brain to represent and process information. It explains the biophysical mechanisms of computation in neurons, computer simulations of neural circuits, and models of learning.

Applications of Computational Neuro Science:

- Deep Learning, Artificial Intelligence and Machine Learning
- Human psychology
- Medical sciences
- Mental models
- Computational anatomy
- Information theory

5.Explain RNN.

Recurrent Neural Networks:

- ❖ RNNs are very powerful, because they combine two properties:

- Distributed hidden state that allows them to store a lot of information about the past efficiently.

- Non-linear dynamics that allows them to update their hidden state in complicated ways.

- ❖ With enough neurons and time, RNNs can compute anything that can be computed by your computer.

Need for RNN:

- ❑ Normal Networks cannot handle sequential data
- ❑ They consider only the current input
- ❑ Normal Neural networks cannot memorize previous inputs

The solution to these issues is the RNN

RNN works on the principle of saving the output of a particular layer and feeding this back to the input in order to predict the output of the layer.

Providing Input to RNN:

We can specify inputs in several ways:

- Specify the initial states of all the units.
- Specify the initial states of a subset of the units.
- Specify the states of the same subset of the units at every time step.

providing Targets to RNN:

We can specify targets in several ways:

- Specify desired final activities of all the units
- Specify desired activities of all units for the last few steps
 - Good for learning attractors
- It is easy to add in extra error derivatives as we backpropagate.
 - Specify the desired activity of a subset of the units