



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-5

2 marks:

1. What is ImageNet, and why is it important in deep learning?

ImageNet is a large-scale dataset containing millions of labeled images used for training deep learning models. It is crucial in computer vision as it led to major breakthroughs in object recognition, especially through the annual ImageNet Large Scale Visual Recognition Challenge (ILSVRC).

2. How does object detection differ from image classification?

Image classification assigns a single label to an entire image, while object detection identifies multiple objects within an image, providing their locations using bounding boxes.

3. What is WaveNet, and how does it generate audio?

WaveNet is a deep generative model developed by DeepMind for speech synthesis. It uses a stack of dilated convolutional layers to generate realistic speech waveforms, capturing temporal dependencies in audio signals.

4. What is the role of Natural Language Processing (NLP) in AI?

NLP enables machines to understand, interpret, and generate human language. It is used in applications like machine translation, sentiment analysis, chatbots, and speech recognition.

5. Explain the working principle of Word2Vec in NLP.

Word2Vec is a neural network-based technique that converts words into continuous vector representations. It uses two models: Continuous Bag of Words (CBOW) and Skip-gram, which learn word relationships based on context.

6. What is Joint Detection, and how is it used in machine learning?

Joint Detection refers to the simultaneous identification and localization of multiple objects in an image or video. It is commonly used in real-time applications like self-driving cars and surveillance.

7. How is deep learning applied in Bioinformatics?

Deep learning is used in Bioinformatics for tasks like protein structure prediction, genomic sequence analysis, drug discovery, and medical diagnostics by identifying patterns in biological data.

8. What are the main challenges in face recognition systems?

Challenges in face recognition include variations in lighting, pose, occlusion, aging effects, and adversarial attacks that can reduce the model's accuracy and reliability.

9. How does scene understanding differ from object detection?

Scene understanding involves recognizing and interpreting the entire scene, including objects, relationships, and context, whereas object detection only identifies and locates individual objects.

10. What is the significance of gathering image captions in computer vision?

Gathering image captions helps in automatic image description generation, aiding accessibility for visually impaired users and improving applications like image search and content understanding.

16 MARKS:

1. Explain the architecture and impact of ImageNet in deep learning. How has it contributed to advancements in computer vision?

Imagenet:

ImageNet is an image database organized according to the WordNet hierarchy (currently only the nouns), in which each node of the hierarchy is depicted by hundreds and thousands of images. In Machine Learning and Deep Neural Networks, machines are trained on a large dataset of various images. Machines are required to learn useful features from these training images. Once learned, they can use these features to classify images and perform many other tasks associated with computer vision. ImageNet gives researchers a common set of images to benchmark their models and algorithms.

ImageNet is useful for many computer vision applications such as object recognition, image classification and object localization. Prior to ImageNet, a researcher wrote one algorithm to identify dogs, another to identify cats, and so on. After training with

ImageNet, the same algorithm could be used to identify different objects. The diversity and size of ImageNet meant that a computer looked at and learned from many variations of the same object. These variations could include camera angles, lighting conditions, and so on. Models built from such extensive training were better at many computer vision tasks. ImageNet convinced researchers those large datasets were important for algorithms and models to work well.

Technical details of Image Net:

ImageNet consists of 14,197,122 images organized into 21,841 subcategories. These subcategories can be considered as sub-trees of 27 high-level categories. Thus, ImageNet is a well-organized hierarchy that makes it useful for supervised machine learning tasks. As many as 1,034,908 images have been annotated with bounding boxes. For example, if an image contains a cat as its main subject, the coordinates of a rectangle that bounds the cat are also published on ImageNet. This makes it useful for computer vision tasks such as object localization and detection. Then there's Scale-Invariant Feature Transform (SIFT) used in computer vision. SIFT helps in detecting local features in an image. ImageNet gives researchers 1000 subcategories with SIFT features covering about 1.2 million images. Images vary in resolution but it's common practice to train deep learning models on sub-sampled images of 256x256 pixels.

ImageNet did not define these subcategories on its own but derived these from WordNet. WordNet is a database of English words linked together by semantic relationships. Words of similar meaning are grouped together into a synonym set, simply called synset. Hypernyms are synsets that are more general. Thus, "organism" is a hypernym of "plant". Hyponyms are synsets that are more specific. Thus, "aquatic" is a hyponym of "plant". This hierarchy makes it useful for computer vision tasks. If the model is not sure about a subcategory, it can simply classify the image higher up the hierarchy where the error probability is less. For example, if model is unsure that it's looking at a rabbit, it can simply classify it as a mammal.

While WordNet has 100K+ synsets, only the nouns have been considered by ImageNet.

How the images are labelled in ImageNet?

In the early stages of the ImageNet project, a quick calculation showed that by employing a few people, they would need 19 years to label the images collected for ImageNet. But in the summer of 2008, researchers came to know about an Amazon service called Mechanical Turk. This meant that image labelling can be crowdsourced via this service. Humans all over the world would label the images for a small fee.

Humans make mistakes and therefore we must have checks in place to overcome them. Each human is given a task of 100 images. In each task, 6 "gold standard" images are placed with known labels. At most 2 errors are allowed on these standard images, otherwise the task has to be restarted.

In addition, the same image is labelled by three different humans. When there's

disagreement, such ambiguous images are resubmitted to another human with tighter quality threshold (only one allowed error on the standard images).

How the images of ImageNet Licensed?

Images for ImageNet were collected from various online sources. ImageNet doesn't own the copyright for any of the images. This has implication on how ImageNet shares the images to researchers.

For public access, ImageNet provides image thumbnails and URLs from where the original images were downloaded. Researchers can use these URLs to download the original images. However, those who wish to use the images for non-commercial or educational purpose, can create an account on ImageNet and request access. This will allow direct download of images from ImageNet. This is useful when the original sources of images are no longer available.

The dataset can be explored via a browser-based user interface. Alternatively, there's also an API. Researchers may want to read the API Documentation. This documentation also shares how to download image features and bounding boxes.

Shortcomings of ImageNet:

Images are not uniformly distributed across subcategories. One research team found that by considering 200 subcategories, they found that the top 11 had 50% of the images, followed by a long tail.

When classifying people, ImageNet uses labels that are racist, misogynist and offensive. People are treated as objects. Their photos have been used without their knowledge. About 5.8% labels are wrong. ImageNet lacks geodiversity. Most of the data represents North America and Europe. China and India are represented in only 1% and 2.1% of the images respectively. This implies that models trained on ImageNet will not work well when applied for the developing world.

2. Describe the working of WaveNet and its applications in speech synthesis. How does it differ from traditional audio generation models?

WaveNet:

WaveNet is a deep generative model of raw audio waveforms. We show that WaveNets are able to generate speech which mimics any human voice and which sounds more natural than the best existing Text-to-Speech systems, reducing the gap with human performance by over 50%. Allowing people to converse with machines is a long-standing dream of human computer interaction. The ability of computers to understand natural speech has been revolutionised in the last few years by the application of deep neural networks. However, generating speech with computers — a process usually referred to as speech synthesis or text to-speech (TTS) — is still largely based on so-called concatenative TTS, where a very large database of short speech fragments are recorded from a single speaker and then recombined to form complete

utterances. This makes it difficult to modify the voice (for example switching to a different speaker, or altering the emphasis or emotion of their speech) without recording a whole new database.

This has led to a great demand for parametric TTS, where all the information required to generate the data is stored in the parameters of the model, and the contents and characteristics of the speech can be controlled via the inputs to the model. So far, however, parametric TTS has tended to sound less natural than concatenative. Existing parametric models typically generate audio signals by passing their outputs through signal processing algorithms known as vocoders. WaveNet changes this paradigm by directly modelling the raw waveform of the audio signal, one sample at a time. As well as yielding more natural-sounding speech, using raw waveforms means that WaveNet can model any kind of audio, including music.

The WaveNet proposes an autoregressive learning with the help of convolutional networks with some tricks. Basically, we have a convolution window sliding on the audio data, and at each step try to predict the next sample value that it did not see yet. In other words, it builds a network that learns the causal relationships between consecutive timesteps (as shown in figure 5.1)

Typically, the speech audio has a sampling rate of 22K or 16K. For few seconds of speech, it means there are more than 100K values for a single data and it is enormous for the network to consume. Hence, we need to restrict the size, preferably to around 8K. At the end, the values are predicted in Q channels (eg. $Q=256$ or 65536), which is compared to the original audio data compressed to Q distinct values. For that, the mulaw quantization could be used: it maps the values to the range of $[0, Q]$. And the loss can be computed either by cross-entropy, or discretized logistic mixture.

The Workflow of WaveNet:

- Input is fed into a causal 1D convolution
- The output is then fed to 2 different dilated 1D convolution layers with sigmoid and tanh activations
- The element-wise multiplication of 2 different activation values results in a skip connection
- And the element-wise addition of a skip connection and output of causal 1D results in the residual

3. Discuss NLP

Natural Language Processing [NLP]:

Language is a method of communication with the help of which we can speak, read and write. For example, we think, we make decisions, plans and more in natural language; precisely, in words. However, the big question that confronts us in this AI era is that can we communicate in a similar manner with computers. In this sense, we can say that Natural Language Processing (NLP) is the sub-field of Computer Science especially Artificial Intelligence (AI) that is concerned about enabling computers to

understand and process human language. Technically, the main task of NLP would be to program computers for analysing and processing huge amount of natural language data.

Natural Language Processing Phases:

The above diagram (Figure5.3) shows the phases or logical steps involved in natural language processing

Morphological Processing

It is the first phase of NLP. The purpose of this phase is to break chunks of language input into sets of tokens corresponding to paragraphs, sentences and words. For example, a word like “uneasy” can be broken into two sub-word tokens as “un-easy”.

Syntax Analysis

It is the second phase of NLP. The purpose of this phase is two folds: to check that a sentence is well formed or not and to break it up into a structure that shows the syntactic relationships between the different words. For example, the sentence like “The school goes to the boy” would be rejected by syntax analyser or parser.

Semantic Analysis

It is the third phase of NLP. The purpose of this phase is to draw exact meaning, or you can say dictionary meaning from the text. The text is checked for meaningfulness. For example, semantic analyser would reject a sentence like “Hot ice-cream”.

Pragmatic Analysis

It is the fourth phase of NLP. Pragmatic analysis simply fits the actual objects/events, which exist in a given context with object references obtained during the last phase (semantic analysis). For example, the sentence “Put the banana in the basket on the shelf” can have two semantic interpretations and pragmatic analyser will choose between these two possibilities.

Content Source: [https:// www.tutorialspoint.com/ natural_language_processing/ natural_language_processing_quick_guide.htm](https://www.tutorialspoint.com/natural_language_processing/natural_language_processing_quick_guide.htm)

Different types based on Working:

1. Speech Recognition — The translation of spoken language into text.
2. Natural Language Understanding (NLU) — The computer’s ability to understand what we say.
3. Natural Language Generation (NLG) — The generation of natural language by a computer.

Applications of NLP:

- ✓ Spam Filters
- ✓ Algorithmic Trading

- ✓ Answering Questions
- ✓ Summarizing Information's etc

4. Explain Word2Vec in detail, including its two main models (CBOW and Skip-gram). How does it help in vector representation of words?

Word2Vec:

Word embedding is one of the most popular representation of document vocabulary. It is capable of capturing context of a word in a document, semantic and syntactic similarity, relation with other words, etc. What are word embeddings exactly? Loosely speaking, they are vector

55

representations of a particular word. Having said this, what follows is how do we generate them? More importantly, how do they capture the context? Word2Vec is one of the most popular technique to learn word embeddings using shallow neural network. It was developed by Tomas Mikolov in 2013 at Google.

The purpose and usefulness of Word2vec is to group the vectors of similar words together in vectorspace. That is, it detects similarities mathematically. Word2vec creates vectors that are distributed numerical representations of word features, features such as the context of individual words. It does so without human intervention.

Given enough data, usage and contexts, Word2vec can make highly accurate guesses about a word's meaning based on past appearances. Those guesses can be used to establish a word's association with other words (e.g. "man" is to "boy" what "woman" is to "girl"), or cluster documents and classify them by topic. Those clusters can form the basis of search, sentiment analysis and recommendations in such diverse fields as scientific research, legal discovery, e-commerce and customer relationship management. Measuring cosine similarity, no similarity is expressed as a 90 degree angle, while total similarity of 1 is a 0 degree angle, complete overlap.

Word2vec is a two-layer neural net that processes text by "vectorizing" words. Its input is a text corpus and its output is a set of vectors: feature vectors that represent words in that corpus. While Word2vec is not a deep neural network, it turns text into a numerical form that deep neural networks can understand.

Word2vec's applications extend beyond parsing sentences in the wild. It can be applied just as well to genes, code, likes, playlists, social media graphs and other verbal or symbolic series in which patterns may be discerned.

Word2vec is similar to an autoencoder, encoding each word in a vector, but rather than training against the input words through reconstruction, as a restricted Boltzmann machine does, word2vec trains words against other words that neighbour them in the input corpus. It does so in one of two ways, either using context to predict a target word (a method known as continuous bag of words, or CBOW), or using a word to predict a target context, which is called skip-gram.

When the feature vector assigned to a word cannot be used to accurately predict

that word's context, the components of the vector are adjusted. Each word's context in the corpus is the *teacher* sending error signals back to adjust the feature vector. The vectors of words judged similar by their context are nudged closer together by adjusting the numbers in the vector.

Similar things and ideas are shown to be "close". Their relative meanings have been translated to measurable distances. Qualities become quantities, and algorithms can do their work. But similarity is just the basis of many associations that Word2vec can learn. For example, it can gauge relations between words of one language, and map them to another.

The main idea of word2Vec is to design a model whose parameters are the word vectors. Then, train the model on a certain objective. At every iteration we run our model, evaluate the errors, and follow an update rule that has some notion of penalizing the model parameters that caused the error. Thus, we learn our word vectors.

5. Describe the Applications of Deep Learning Networks: Joint Detection

Applications of Deep Learning Networks: Joint Detection

Deep Learning finds a lot of usefulness in the field of Biomedical and Bioinformatics. Deep Learning algorithms can be used for detecting fractures or anatomical changes in the human bones or in bone joints thereby early prediction of various diseases like arthritis can be done which helps in early curing of the so-called diseases also. Knee osteoarthritis (OA) is a very general joint disease that disturbs many people especially people over 60. The severity of pain caused by knee OA is the most important portent to disable. Until now, the bad impact of osteoarthritis on health care and public health systems is still increasing.

The Convolution Filter is made to move over the image. The stride value here considered for this case study is 1. We have used Max Pooling and in the fully connected layer we have used Softmax aggregator.

The above diagram is a schematic representation of a Deep Learning based network which can be used for human knee Joint deformities Identification purpose.

The Convolution filter moves to the right with a certain Stride Value till it parses the complete width. Moving on, it hops down to the beginning (left) of the image with the same Stride Value and repeats the process until the entire image is traversed. The Kernel has the same depth as that of the input image. The objective of the Convolution Operation is to extract the high-level features such as edges, from the input image. Stride is the number of pixels shifts over the input matrix. When the stride is 1 then we move the filters to 1 pixel at a time. When the stride is 2 then we move the filters to 2 pixels at a time and so on

Pooling layers section would reduce the number of parameters when the images are too large. Spatial pooling also called subsampling or down sampling which reduces the dimensionality of each map but retains important information. This

is to decrease the computational power required to process the data by reducing the dimensions

Types of Pooling:

- Max Pooling
- Average Pooling
- Sum Pooling
- The image is flattened into a column vector.
- The flattened output is fed to a feed-forward neural network and backpropagation applied to every iteration of training.

Over a series of epochs, the model is able to distinguish between dominating and certain low-level features in images and classify them using the Softmax Classification technique. The feature map matrix will be converted as vector (x_1, x_2, x_3, \dots). These features are combined together to create a model.

59

Finally, an activation function such as softmax or sigmoid is used to classify the outputs as Normal and Abnormal.

Steps Involved:

- Provide input image into convolution layer
- Choose parameters, apply filters with strides, padding if requires. Perform convolution on the image and apply ReLU activation to the matrix.
- Perform pooling to reduce dimensionality size
- Add as many convolutional layers until satisfied
- Flatten the output and feed into a fully connected layer (FC Layer)
- Output the class using an activation function (Logistic Regression with cost functions) and classifies images.

Other Applications:

Similarly for the other Applications such as Facial Recognition and Scene Matching applications appropriate Deep Learning Based Algorithms such as AlexNet, VGG, Inception, ResNet and or Deep learning-based LSTM or RNN can be used. These Networks has to be explained with necessary Diagrams and appropriate Explanations.