Fundamentals of Artificial Intelligence
Prof. Shyamanta M. Hazarika
Department of Mechanical Engineering
Indian Institute of Technology-Guwahati

Lecture - 34
Deep Learning: A Brief Overview

Welcome to Fundamentals of Artificial Intelligence. We continue our discussion on machine learning. In the last lecture, we had seen how the human brain inspired simple computational elements and led to the development of artificial neural networks. Starting with McCulloch and Pitts model, the artificial neural networks in its beginning caused quite is stir.

However, by the 60s interest in artificial neural networks somehow waned out and it was only in the 90s that people again started looking at the artificial neural networks as a computational paradigm. Over the last decade, however, there has been a huge surge in the use of a variant of artificial neural networks which have a number of layers in it referred to as deep neural networks within artificial intelligence. Today we look at deep learning.

(Refer Slide Time: 02:21)

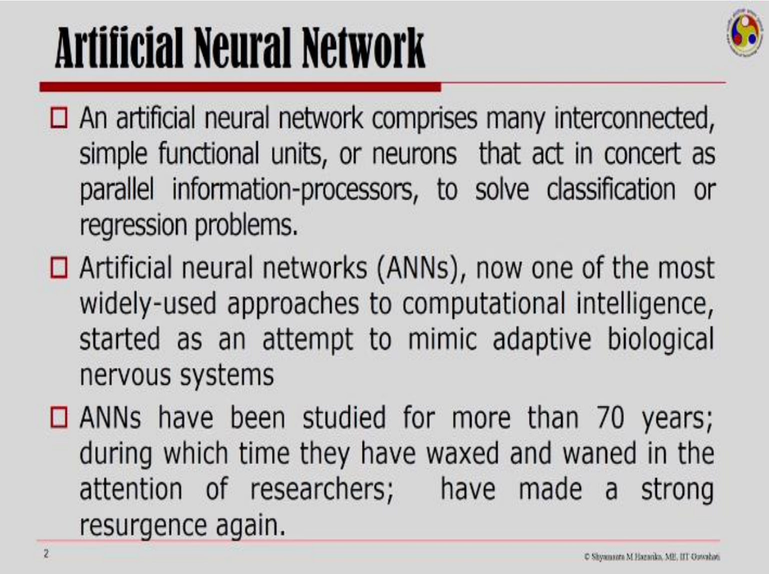

Now deep learning has its roots in the artificial neural network and the McCulloch and Pitts model of the neuron. An artificial neural network if you recall comprises many interconnected simple functional units or what is referred to as neurons that act in concert or parallel to solve classification or regression problems. Recall from our last lecture that individually if you look at the computational power of the computer is hugely more than what you get in the human brain in terms of a single neuron.

However, the brain functions with all the neurons working in unison or in concert and that gives rise to tremendous power which is not achievable by a computational system. Artificial neural networks aim to create this structure of neurons and now comprises one of the most widely used approaches to computational intelligence, which started as an attempt to mimic adaptive biological nervous systems.

Artificial neural networks has been studied for more than 70 years. And as I was mentioning in the beginning of this lecture, during this time artificial neural networks have waxed and waned attention of researchers and only in the last decade, they have made a strong resurgence again.

(Refer Slide Time: 04:26)

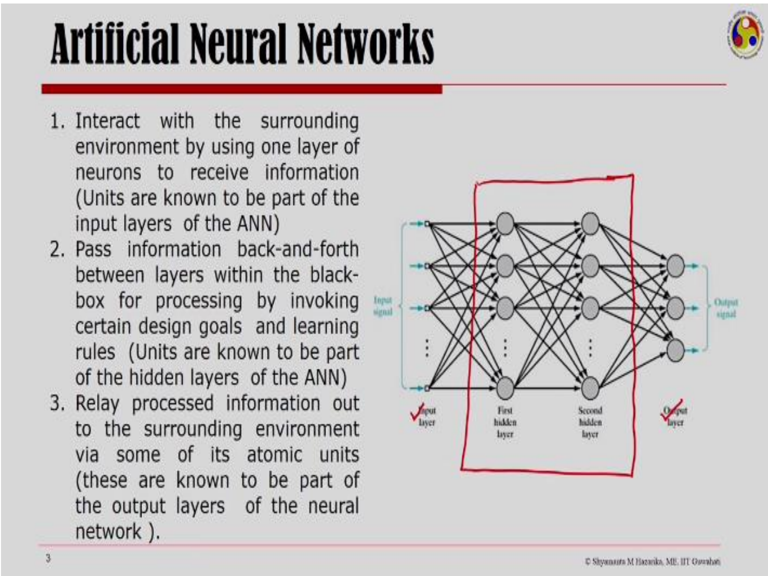

Artificial neural networks as the name suggest, is a network of artificial neurons. And we saw the very simplest of which is the perceptron and in the last lecture, we had looked at how a perceptron can lead to a multiple layer neural network. Now what we have here on your screen is a multilayer artificial neural network. A multilayer artificial neural network interacts with the surrounding environment by using one layer of neurons to receive information.

So this layer is referred to as the input layer of the artificial neural network. This input layer pass information on to the next layer. In fact, information is passed back and forth between the layers for processing by invoking certain design goals and certain learning rules. There can be a number of layers in between the input and the final layer. So these layers that we have between the input and the final layer are referred to as the hidden layers.

The hidden layers pass information and as we will see, they allow adjustment of weights so that the process information that is relayed out to the surrounding environment from the final layer referred to as the output layer reflects some amount of learning.

(Refer Slide Time: 06:39)

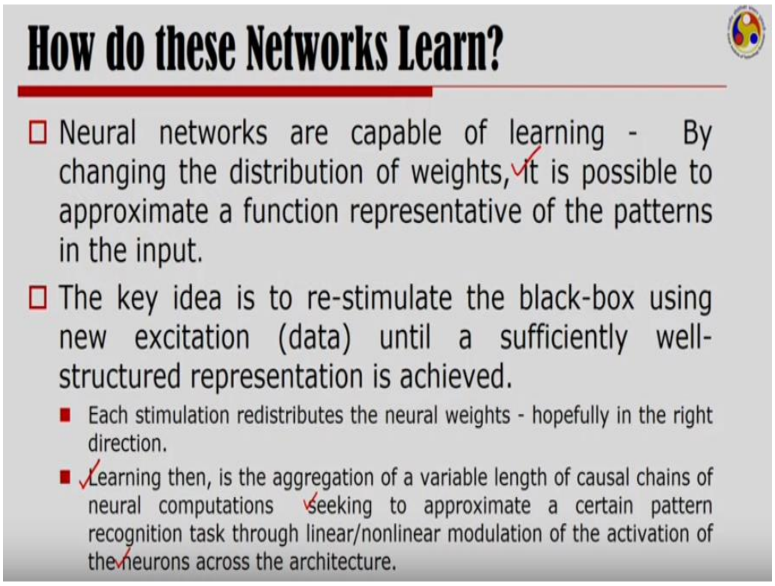

Now the question is how do these networks learn? Neural networks are capable of learning by changing the distribution of weights. So it is possible to approximate a function which is a representative of the patterns in the input. And the key idea is to re-stimulate the black box using new excitation until a sufficiently well structured representation is achieved.

And this is done by the hidden layers between the input and the final output layer. Each stimulation actually redistributes the neural weights hopefully in the right direction and learning actually can be seen in neural networks as the aggregation of a variable length of causal chains of these neural computations, which seek to approximate a certain pattern recognition task to either linear or nonlinear modulation of the activations of all these neurons, if we allow them to be called as across the architecture.

(Refer Slide Time: 08:14)

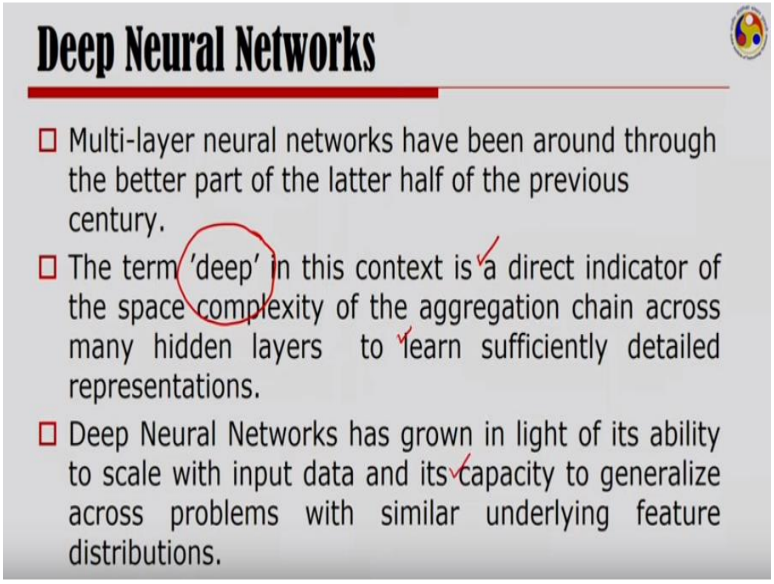

Now, we should realize that multi-layer neural networks have been around the better part of the later half of the previous century. The term deep that is used now is in the context a direct indicator of the space complexity of the aggregation chain across this many hidden layers that I was referring to, to learn sufficiently detailed representation.

So the very term deep here that we now use in the context of artificial neural networks is about what is the space complexity of the aggregation chains across the many hidden layers. So roughly put, it is a reflection of the number of hidden layers that you have in an artificial neural network. So deep neural networks has taken the concept of artificial neural network through a great resurgence and has grown in light of its ability to scale with input data and its capability to generalize across problems with similar underlying feature distributions.

(Refer Slide Time: 09:46)

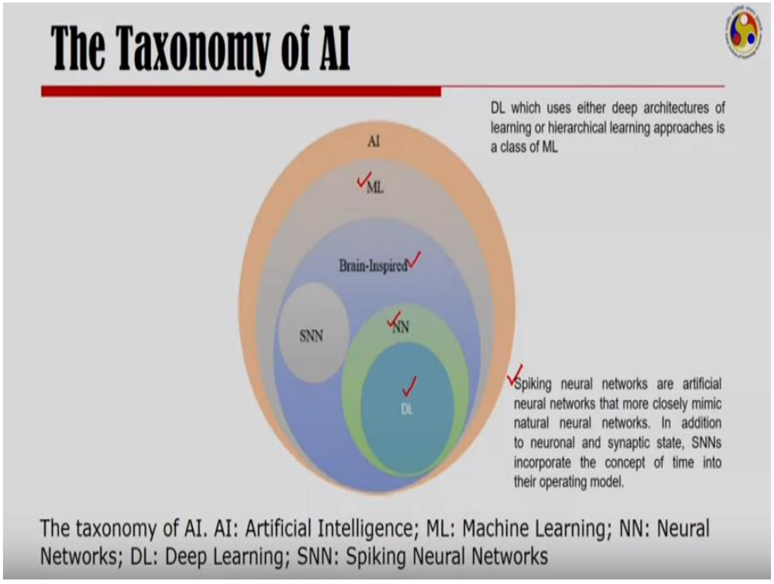

If you look at the taxonomy of artificial intelligence, you will see that I have been emphasizing this in a number of lectures that AI comprises of the very idea of machine learning. So machine learning can be seen as a subset of AI and within machine learning, we have deep learning which is an idea to use deep architectures of learning or hierarchical learning approaches within machine learning using artificial neural networks.

So within AI, we have machine learning, which is a subset of AI and with the machine learning, the brain inspired architecture of computation, the connectionist AI or the connectionist way of doing computation, which is artificial neural nets have been made with the idea of deep learning. Another closely related idea that needs mention here within brain inspired artificial neural networks is the idea of spiking neural networks.

Spiking neural networks are artificial neural networks that more closely mimic the natural neural networks. So in addition to the neuronal and the synaptic state that is what the artificial neural network tries to mimic. The spiking neural networks also incorporate the concept of time into their operating models. So the very idea of deep neural networks is to use deep architectures of learning or hierarchical learning approaches.

(Refer Slide Time: 11:58)

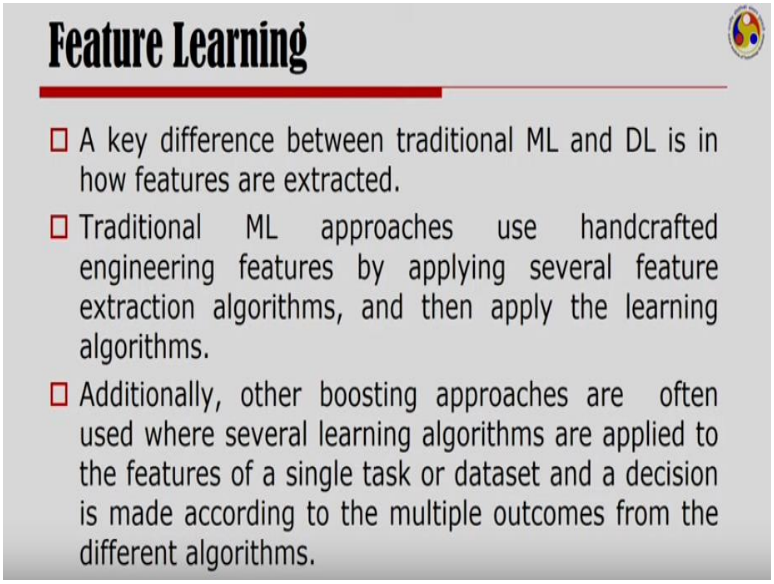

Now one very interesting thing that needs to be taken note of at this point is the idea of feature learning that makes traditional machine learning and deep learning different. So the differences between how features are extracted in traditional machine learning approaches, we use handcrafted engineering features by applying several feature extraction algorithms and then apply the learning algorithms on them.

Whereas, in deep learning, this distinction is not there. In traditional machine learning, additionally we use boosting approaches where several learning algorithms are applied to the features of a single task or data set and a decision is made according to the multiple outcomes from the different algorithm.

(Refer Slide Time: 13:01)

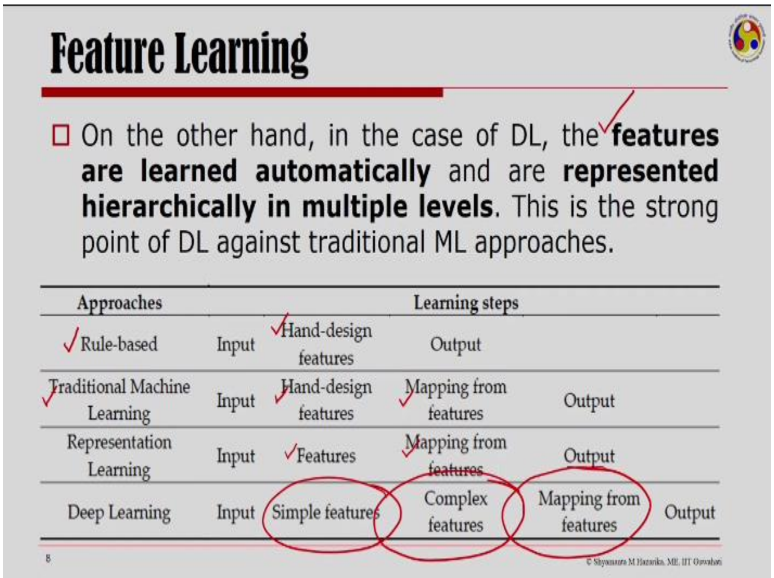

Whereas in deep learning, the features are automatically learned and they are represented hierarchically in multiple levels. So if you look at certain traditional machine learning approaches like a rule based system, here, we would have hand-design features and then we will get to the output. Traditional machine learning will also look at hand-design features.

Thereafter, learning involves mapping from features finally to the output. Representation learning is about getting to the features and mapping from the features to the output. Whereas what makes deep learning interesting is that the very idea of having features and mapping from the features all of these huge process is done within the artificial neural network itself.

This whole idea of having features and then extracting complex features from them and mapping from features to outputs is all happening within the artificial neural network structure.

(Refer Slide Time: 14:37)

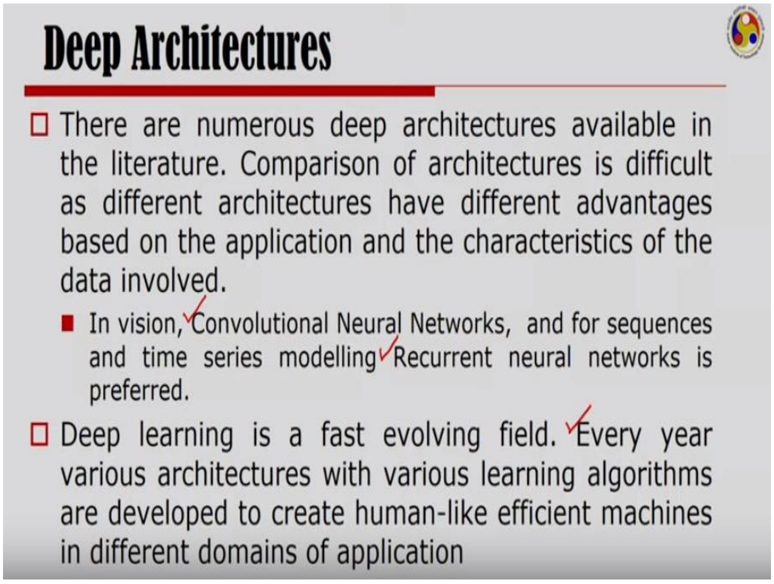

There are numerous deep architectures available in the literature. Now comparison of architectures is difficult as different architectures have different advantages based on the application and the characteristics of data involved, like in vision, convolutional neural networks, and for sequence and time series modeling, recurrent neural networks are the preferred choices.

Deep learning however is a fast evolving field. And therefore, as part of a fundamental course in AI, it is pertinent to have a look at a very basic idea of what is deep learning and what are the different architectures involved. Here, it is neither intended nor feasible, in a short duration of a single lecture to cover all the details of deep learning architectures available.

This lecture is rather to be looked at as an exploratory exposure, where we look at the fundamental architectures being used and emphasize on the concept rather than on the mathematical regard. Now the very idea of looking at architectures of deep learning is also difficult by the fact that every year various architectures with various learning algorithms are developed to create human like efficient machines in different domains of application.

So what I intend to do in the remainder of the lecture is to give you some idea of the different architectures involved in deep neural networks.

(Refer Slide Time: 16:46)

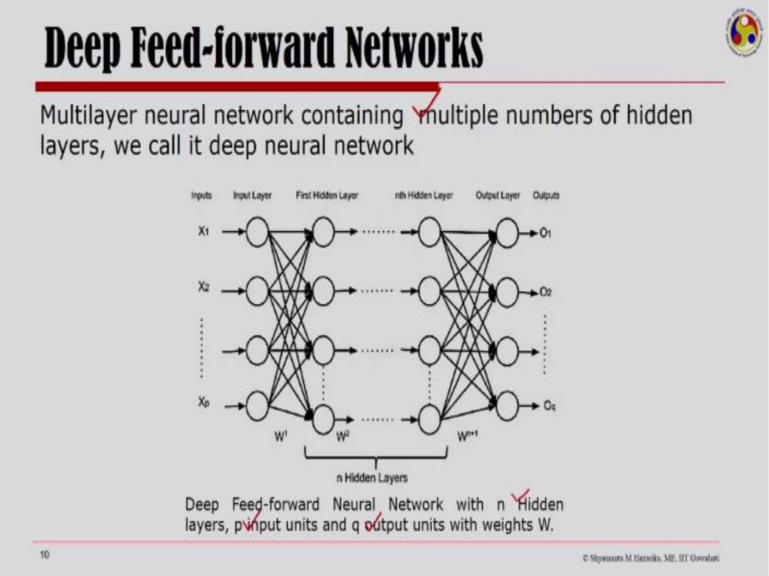

One of the basic one is referred to as the deep feed-forward network. So a deep feed-forward network has a number of hidden layers and a number of input units leading to a number of output units. Now recall that in the last lecture, we had seen a multilayer perceptron network. So a deep feed-forward network is actually a multilayered neural network, but then it contains multiple number of layers. And this huge number of layers is what makes it a deep neural network.

(Refer Slide Time: 17:37)

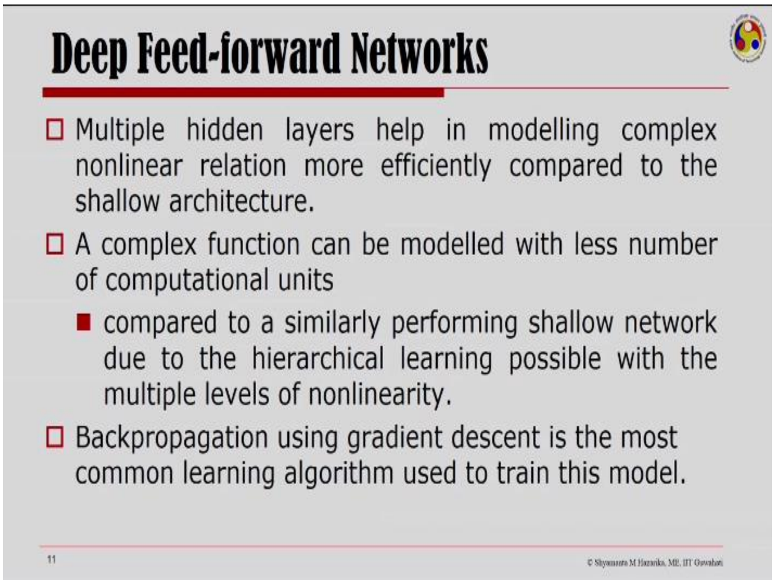

Now the multiple hidden layers, help in modeling complex nonlinear relation more efficiently compared to the shallow architecture. A complex function can be modeled with less number of computational units compared to a similarly performing shallow network due to the hierarchical learning possible with the multiple levels of non linearity that is introduced in a deep feed-forward network.

Back propagation using gradient descent is the most common learning algorithm used to train this model.

(Refer Slide Time: 18:20)

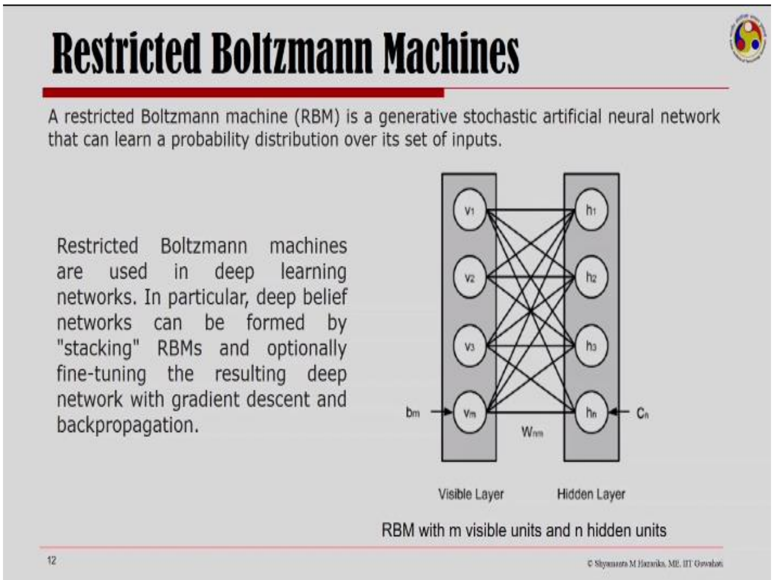

The next off the architectures is the restricted Boltzmann machine. A restricted Boltzmann machine is a generative stochastic artificial neural network that can learn a probability distribution over a set of inputs. So here we have a restricted Boltzmann machine. The restricted Boltzmann machine are used in deep learning networks, in particular, in deep belief networks.

And these deep belief networks as we will see in a short while, can be formed by stacking this restricted Boltzmann machines and fine-tuning the resulting deep network with gradient descent and back propagation.

(Refer Slide Time: 19:14)

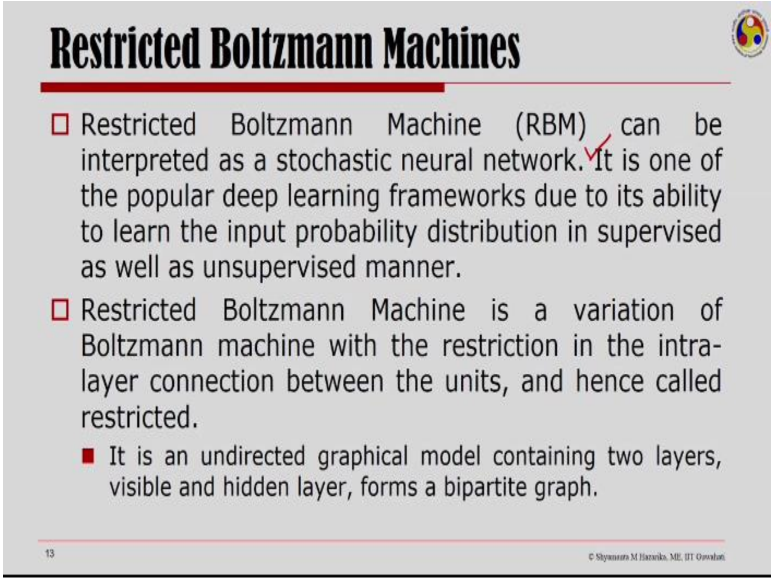

Restricted Boltzmann machine actually can be interpreted as a stochastic neural network. It is one of the popular deep learning framework due to its ability to learn the input probability distribution in supervised as well as unsupervised manner. Restricted Boltzmann machine is a variation of a Boltzmann machine with the restriction in the interlayer connection between the units and hence it is called restricted.

Now it is an undirected graphical model, where we have two layers as seen in the previous slide, a visible and a hidden layer and it forms a bipartite graph.

(Refer Slide Time: 20:07)

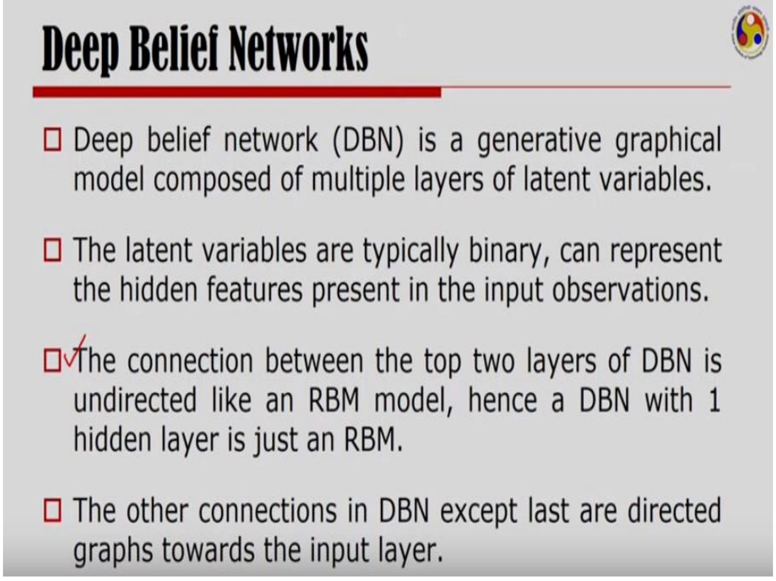

Deep belief networks is a generative graphical model. Now they are composed of multiple layers of latent variables. The latent variables are typically binary, they can represent the hidden features present in the input observation. Now as I was telling you when discussing the restricted Boltzmann machine, the connection between the two top layers of the deep belief network is undirected like a restricted Boltzmann machine.

Hence a deep belief network with one hidden layer is just a restricted Boltzmann machine. The other connections in a deep belief network except the last are directed graph towards the input layers.

(Refer Slide Time: 21:01)

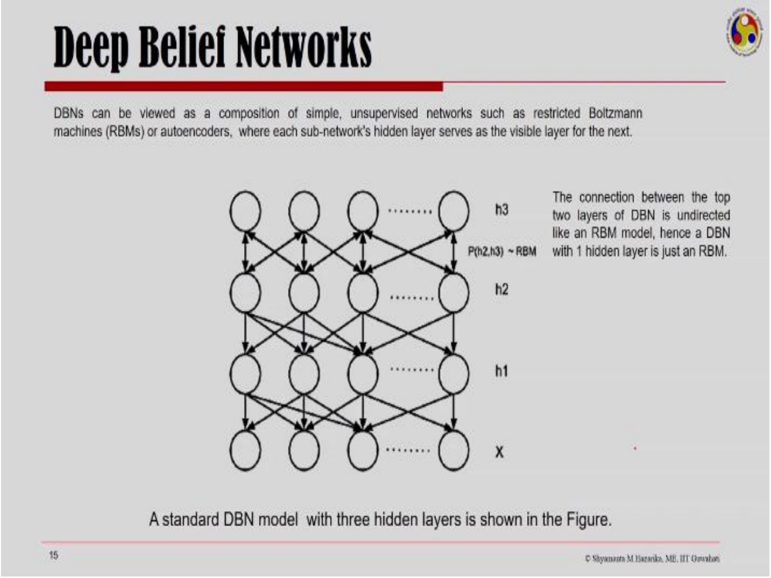

Now here on your screen is a standard deep belief network model with three hidden layers. And this can be viewed as a composition of simple unsupervised network, where each sub network's hidden layer serve as the visible layer for the next one. And the connection between the two layers as already mentioned, at this top portion is undirected like a restricted Boltzmann machine.

(Refer Slide Time: 21:38)

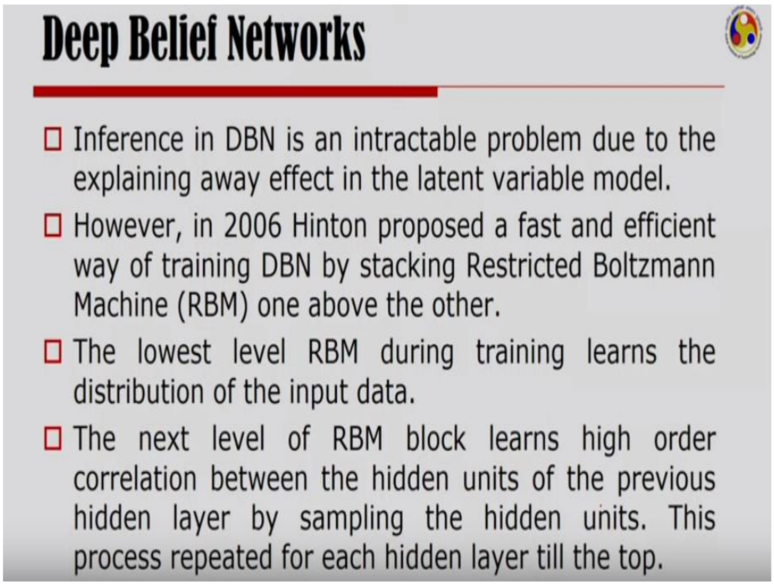

Inference in a deep belief network is an intractable problem due to the explaining away effect in the latent variable model. However, in 2006 Hinton proposed a fast and efficient way of training the deep belief networks by stacking restricted Boltzmann machine one above the other and this is when artificial neural networks took a huge turn and there was resurgence in the use of ANNs, particularly deep neural networks within artificial intelligence.

So even if the pre history of artificial neural networks date backs to the McCulloch Pitts model of the 40s followed by Rosenblatt's perceptron in the 50s. It was only in 2006, with Hinton's proposal of a fast and efficient way of training deep belief networks that changed the very course of deep neural networks.

Now in this idea of stacking restricted Boltzmann machine one above the other, the lowest level restricted Boltzmann machine during training, learns the distribution of the input data. The next level of the RBM block learns higher order correlation between the hidden units of the previous hidden layer by sampling the hidden units and this is repeated for each hidden layer till the top.

(Refer Slide Time: 23:23)

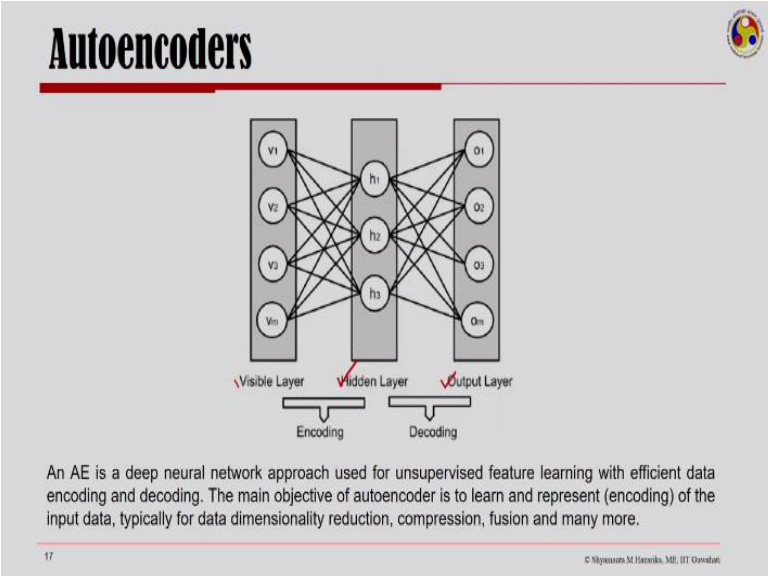

The next sort of deep architectures is what I referred to as autoencoders. An autoencoder is a deep neural network, which is used for unsupervised feature learning with efficient data encoding and decoding. So the main objective of autoencoder is to learn and represent input data. Typically, this is used for operations such as dimensionality reduction, compression, fusion, and many other similar operations.

So what we have in an autoencoder is a visible layer followed by a hidden layer and then an output layer.

(Refer Slide Time: 24:18)

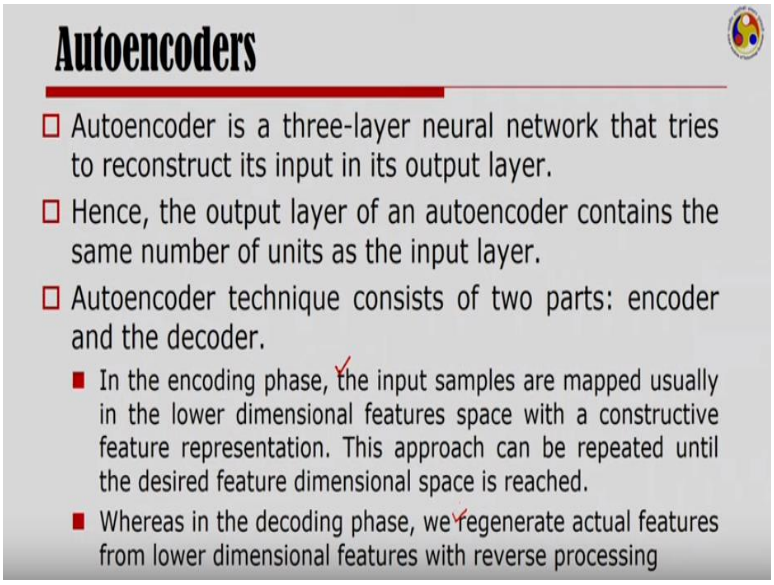

So it is a three layer neural network that tries to reconstruct its input in its output layer. The output layer of an autoencoder, therefore contains the same number of units as the input layer. The autoencoder technique consists of two parts as I was referring to. You have the encoder, which is between the input and the hidden layer and you have the decoder which is between the hidden and the output layer.

So in the encoding phase, the input samples are mapped usually in the lower dimensional feature space, we take constructive feature representation. And this approach can be repeated until you have the desired feature dimensional space. Whereas in the decoding phase, we regenerate the actual features from the lower dimensional features with reverse processing.

And now therefore you could see how autoencoders could be used for tasks such as dimensionality reduction.

(Refer Slide Time: 25:36)

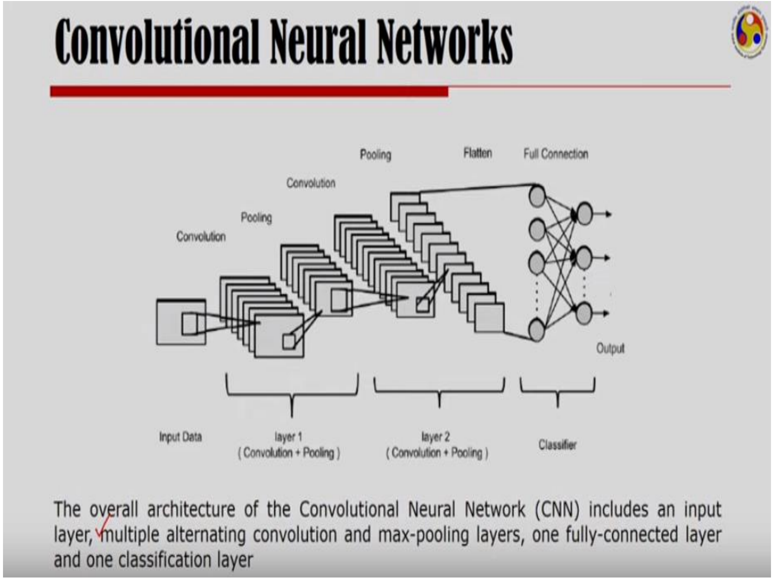

The next off the deep architectures is the most well known convolutional neural networks. Now the convolutional neural networks contains an input layer and it contains multiple alternating convolutional and maximum-pooling layers. And finally, a fully connected layer and then a classification layer.

So the layers involved in any convolutional neural network model are the convolutional layers and the sub sampling or the pooling layers, which allow the network to learn filters that are specific to specific parts in the data. So convolutional neural networks are mostly used for image processing and therefore the network learn the filters that are specific to specific parts in an image.

(Refer Slide Time: 26:51)

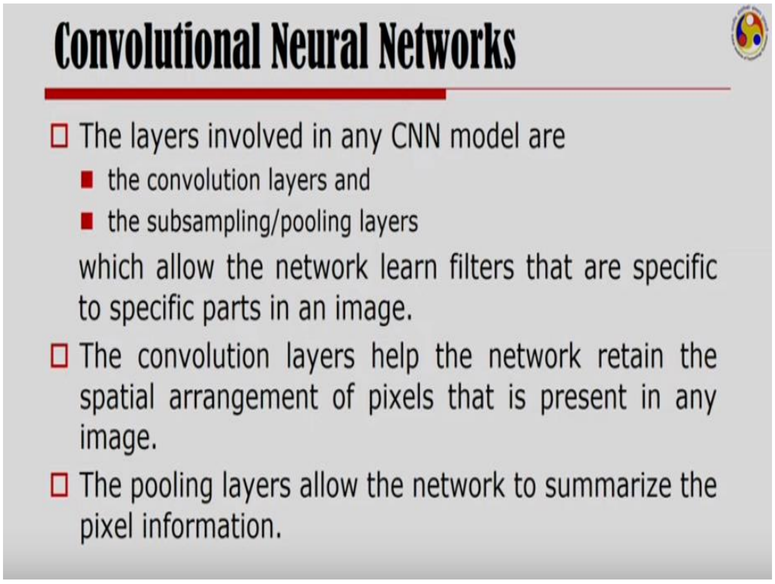

The conventional layers help the network retain the special arrangement of pixels that is present in any image. The pulling layers on the other hand, allow the network to summarize the pixel information. So basically, you have now an image that has spatial information, because of the spatial arrangement of pixels. And the convolutional layer helps to retain the spatial information or the spatial arrangement.

And the pooling layer allows the network to summarize the pixel information. Convolutional neural networks is possibly one of the largest supplied of the deep learning architectures. Now what I have presented here is a very simple architecture, which basically covers the pipeline of the convolutional neural network.

There are different variants of the convolutional neural network based on how you do the convolution layers or how you do the pooling, giving rise to altogether different types of convolutional neural networks, which we have not covered.

(Refer Slide Time: 28:21)

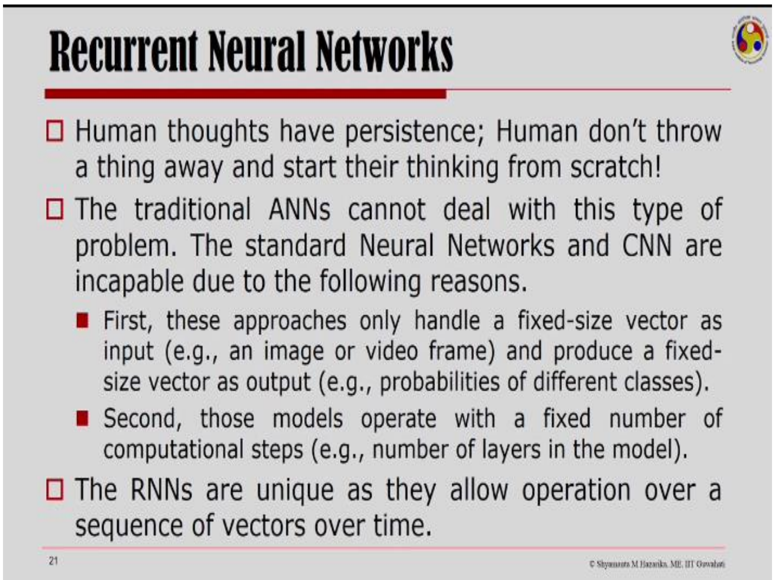

The next deep learning architecture that we will cover is the recurrent neural networks. So this is based on the very idea that human thoughts have persistence. The traditional artificial neural networks cannot deal with this type of problem. The standard neural networks, including the convolutional neural networks are incapable of doing this due to the following reasons.

One, these approaches can only handle a fixed-size vector as input, example an image or a video frame and it produces a fixed size vector as output. That is the probabilities of different classes for example. Number two, these models operate with a fixed number of computational steps. That is, have a fixed number of layers in the model. The recurrent neural networks are unique in the sense that they allow operation over a sequence of vectors over time.

(Refer Slide Time: 29:35)

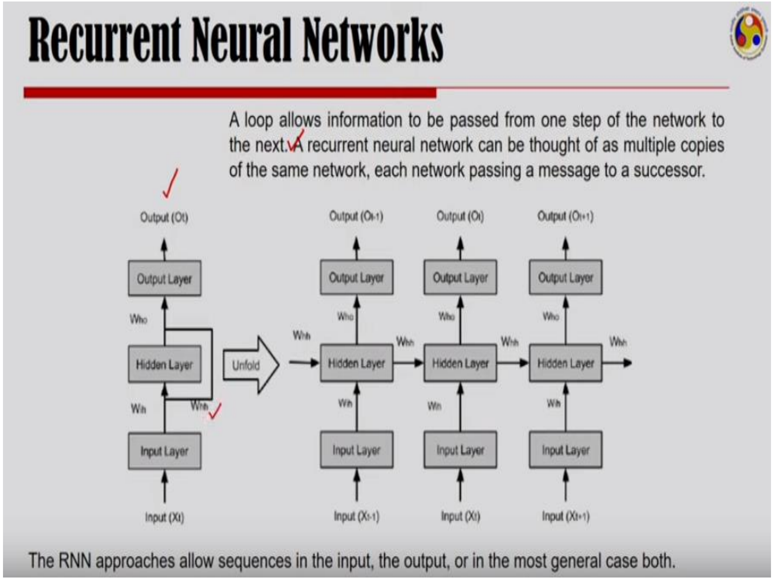

So here on your screen, we have the recurrent neural network approach, where we show this recurrent neural network on the left and the recurrent neural network can be seen as unfolding to allow multiple copies of the same network, each network passing a message to the successor. So a recurrent neural network allows sequences in the input and output or in the most general cases, both of them.

And a recurrent neural network can be thought of as multiple copies of the same network. In fact, in recurrent neural network, a loop allows the message or the information to be passed from one step of the network to the next.

(Refer Slide Time: 30:34)

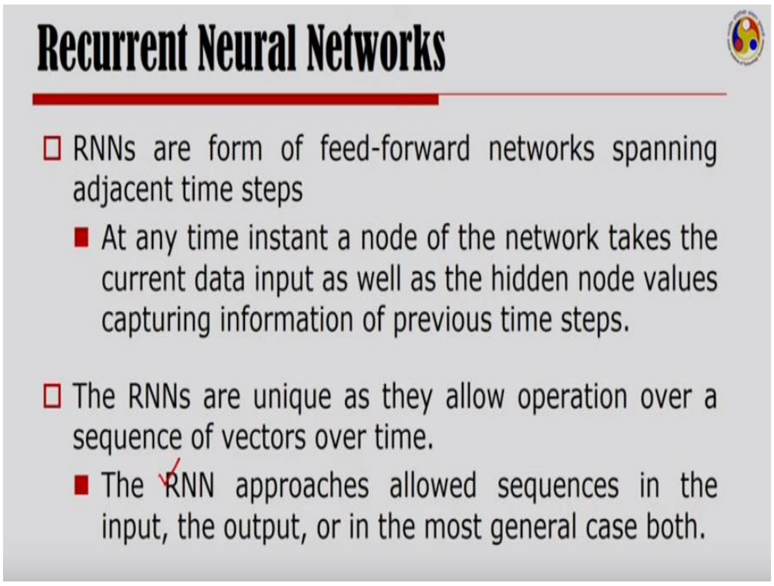

Recurrent neural networks are form of feed-forward networks spanning adjacent time steps. At any one time instant a node of the network takes the current data input as well as the hidden node values, capturing information of previous time steps. And the recurrent neural network approach allows sequence in the input, output or in the most general case, in both of them as already mentioned.

(Refer Slide Time: 31:06)

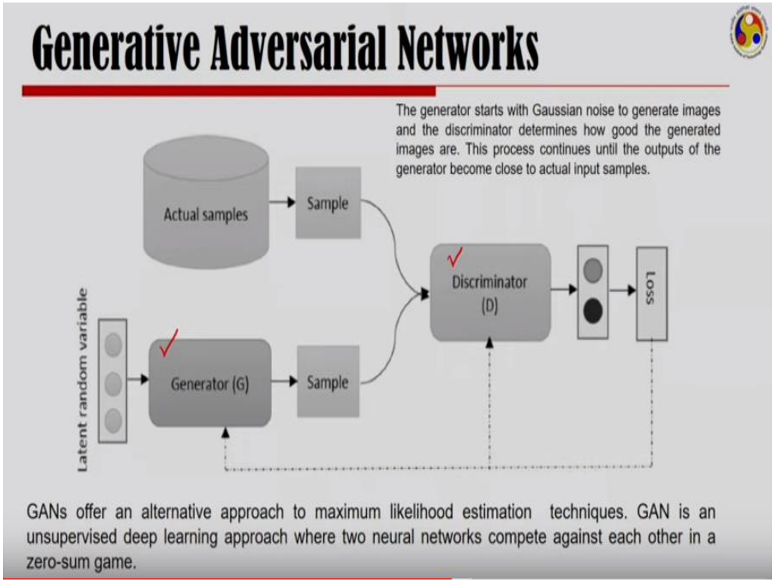

The other deep architecture is the generative adversarial networks. Now these generative adversarial networks offer an alternative approach to maximum likelihood estimation techniques and is an unsupervised deep learning approach where two neural networks compete against each other in a zero-sum game. So basically, the generator starts with possibly Gaussian noise to generate images.

And the discriminator determines how good the generated images are. And this process continues, until the output of the generator becomes close to the actual input samples.

(Refer Slide Time: 32:01)

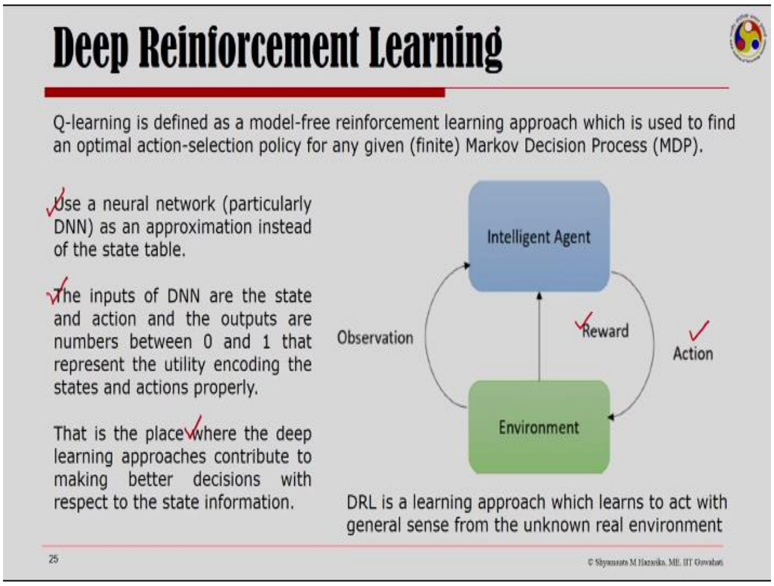

Finally, let us take a look at what is referred to as the deep reinforcement learning. Recall that in our lecture on reinforcement learning, we had seen that a reinforcement learning approach is where the agent learns to take an action based on the reward. And we have looked at an algorithm called the Q-learning. So the Q-learning is a modern free reinforcement learning approach, which is used to find an optimal action selection policy for any given Markov decision process.

Now here in deep reinforcement learning, you use a neural network, particularly a deep neural network as an approximation instead of the state table. And the inputs of the DNN are the state and action and the outputs are number between 0 and 1 that represents in a way the utility encoding that the states and actions. And this is where the deep learning approach contributes to making better decisions with respect to the state information.

(Refer Slide Time: 33:36)

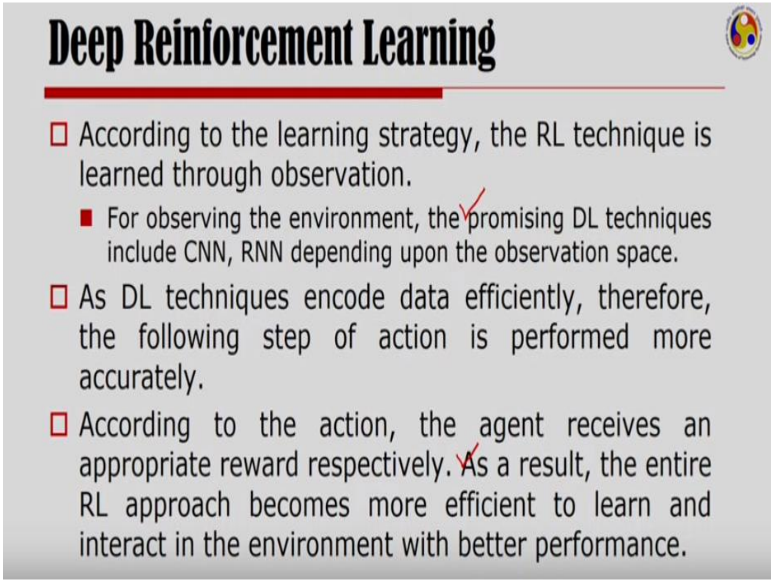

Deep reinforcement learning is about observing the environment and using promising deep learning techniques including convolutional neural networks or recurrent neural networks depending on the observation space. Now as deep learning techniques encode data efficiently, the idea is that the following step of action would be performed more accurately.

Now in reinforcement learning according to the action, the agent receives an appropriate reward. And as a result, the entire reinforcement learning approach becomes more efficient to learn and interact in the environment with better performance using deep learning techniques, instead of just a table look up.

(Refer Slide Time: 34:37)

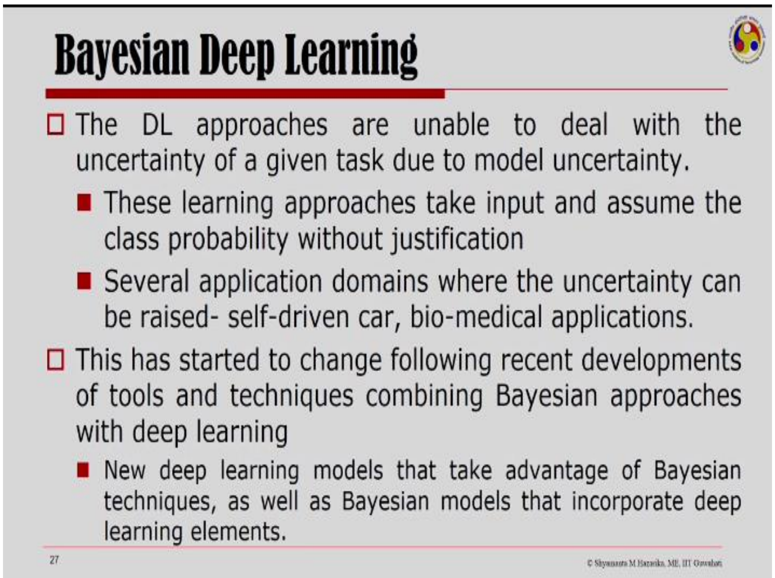

The deep learning approaches that we have been looking at are in fact unable to deal with uncertainty of a given task due to model uncertainty. These learning approaches take input and assume the class probability without justification. Whereas several application domains are there where uncertainty can be raised like domains of self driven car or biomedical applications.

This has led to recent developments of tools and techniques that combine Bayesian approaches with deep learning. New deep learning models that take advantage of Bayesian techniques as well as Bayesian models that incorporate deep learning elements are being looked into.

(Refer Slide Time: 35:36)

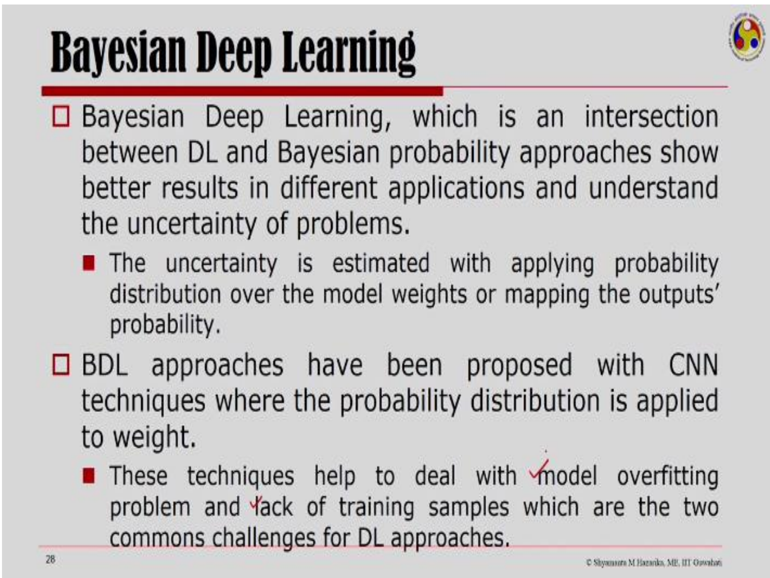

Bayesian deep learning therefore is an intersection between deep learning and Bayesian probability approaches. And as you can see, have promising results in different applications, where uncertainty is involved. The uncertainty is estimated with applying probability distribution over the model weights or mapping the outputs probability.

Bayesian deep learning approaches have been proposed with convolutional neural network techniques. These techniques help to deal with model overfitting problem, and even lack of training samples, which are two common challenges for deep learning approaches.

(Refer Slide Time: 36:32)

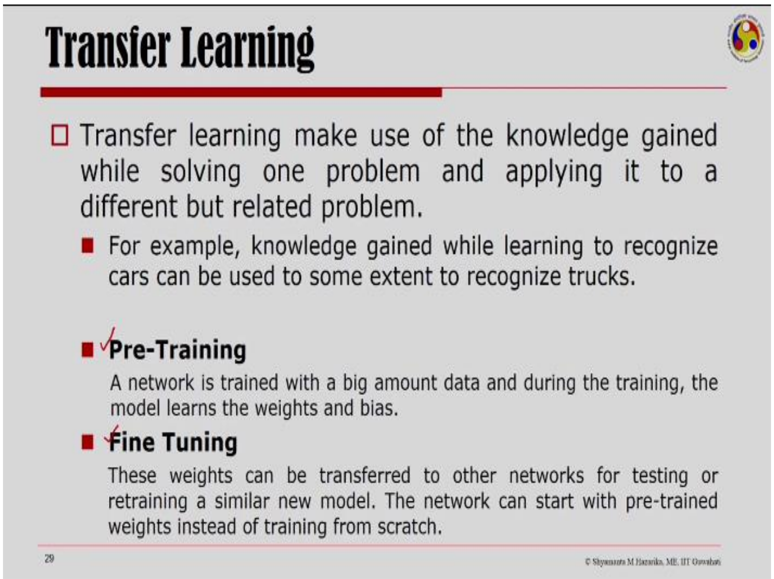

Now transfer learning is an idea within deep learning, which is making use of the knowledge that is gained while solving one problem and applying it to a different but related problem. For example, the knowledge gained while learning to recognize cars can be used to some extent, to recognize trucks. And this is a very interesting development within deep neural networks.

So here, two phases are involved. You have pre-training, where the idea is to train a network with big amount of data so as the model learns the weights and the bias, and then do fine-tuning, which is that these weights are transferred to other networks for testing or training a similar new model. And the network can start with the pre-trained weights, instead of training from scratch.

(Refer Slide Time: 37:40)

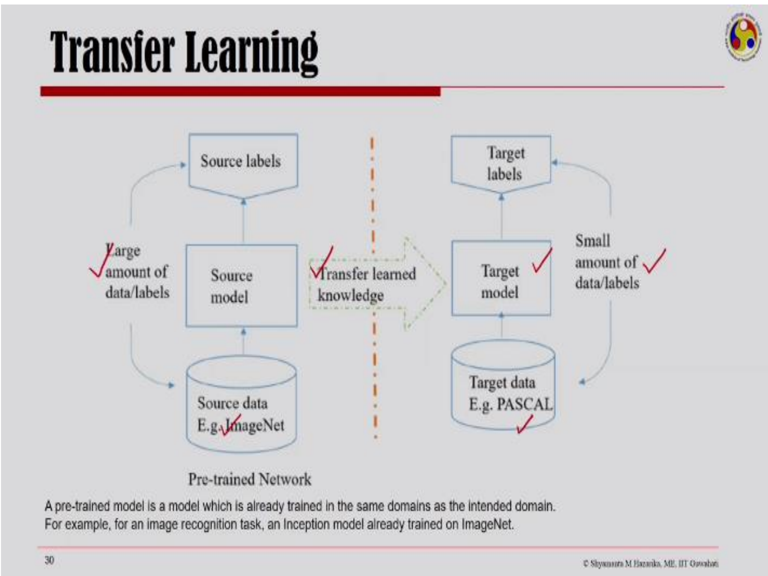

Now here on your screens, I have just shown a pre-trained model, which is already trained over the ImageNet, and then you transfer the learned knowledge to your target model, and then try to use it over a new data set. Now the idea is that you would have small amount of data and labels working for you on the transferred model, whereas you would have large amount of data and labels when you are pre-training.

(Refer Slide Time: 38:20)

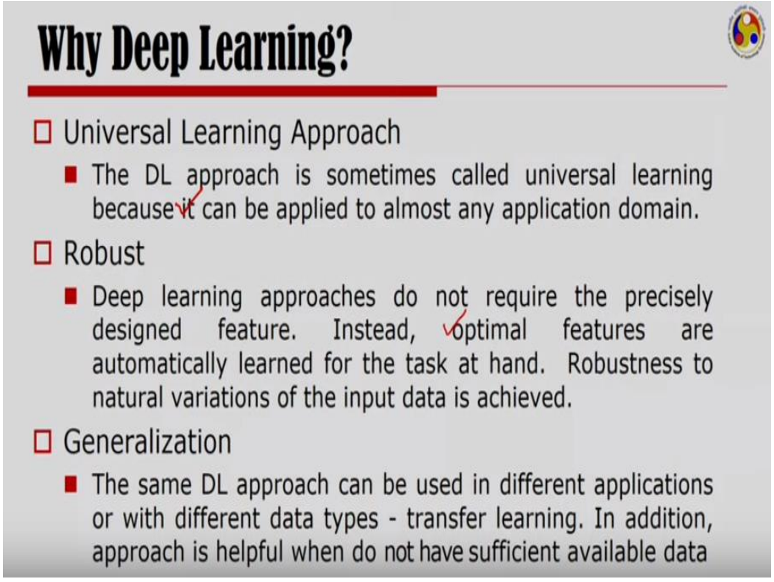

Having looked at very quickly, these deep learning architectures and certain ideas within deep learning that is emerging, a question that we have not looked at is why do we prefer deep learning? Now the first and foremost is that the deep learning approach is a universal learning approach because it can be applied to almost any application domain.

Number two, the deep learning approach is robust. Deep learning approaches do not require the precisely designed features. Instead, optimal features are actually automatically learned for the task at hand and robustness to natural variations of the input data is achieved.

Generalization is something that makes deep learning powerful. The same deep learning approach can be used in different applications or with different data types. And that is exactly what we saw in the last slide on transfer learning. In addition, the approach is helpful when we do not have sufficient data. So generalization makes deep learning powerful.

(Refer Slide Time: 39:57)

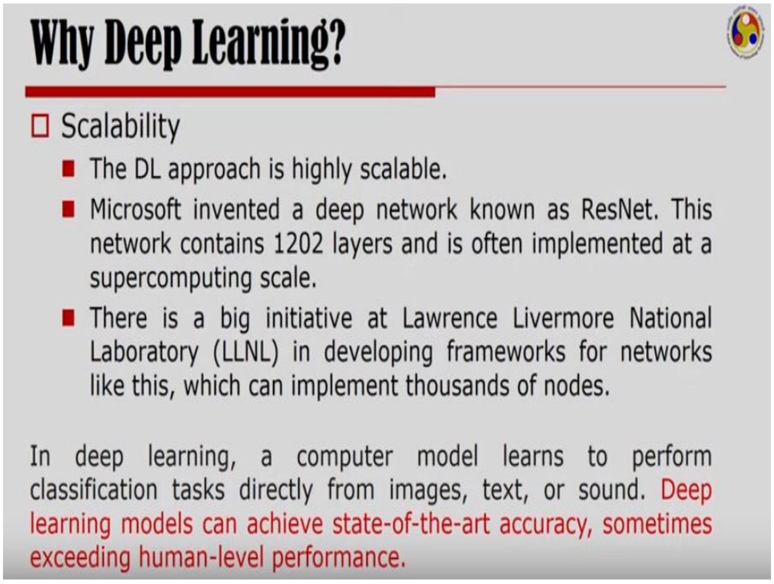

And of course, its winning point is scalability. The deep learning approach is highly scalable. In fact, Microsoft invented a deep network, known as ResNet. And this network contains 1202 layers and is often implemented at a supercomputing scale. Now there is a big initiative at developing frameworks for networks like this, which can implement thousands of nodes.

In deep learning, a computer model learns to perform classification tasks directly from images, text, or sound. And we do not need any human crafted feature extraction processes. Deep learning methods can achieve state of the art accuracy, sometimes even exceeding human level performance. And this is what makes deep learning so exciting.

(Refer Slide Time: 41:03)

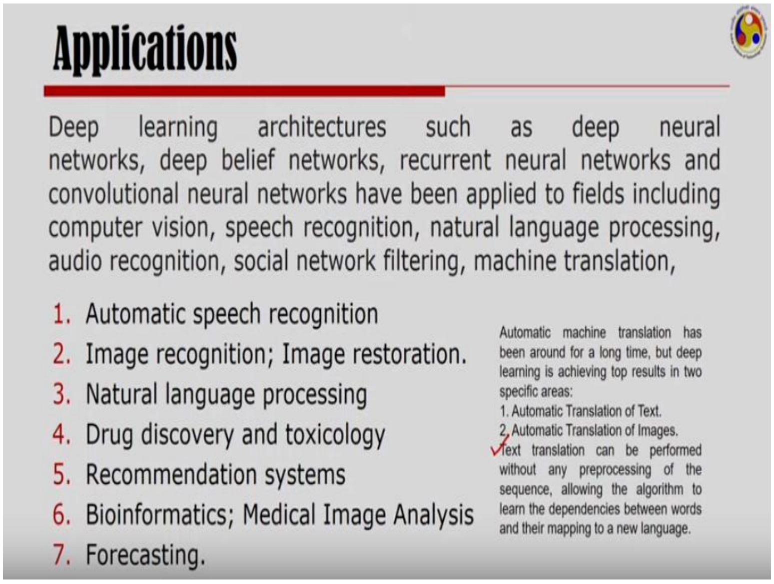

Now deep learning architectures, such as deep neural networks, deep belief networks, recurrent neural networks, convolutional neural networks, have been applied to fields including computer vision, speech recognition, natural language processing, audio recognition, social network filtering, and machine translation. Here, I have listed only seven of them. But then there are numerous others.

Automatic machine translation has been around for a long time. But deep learning is achieving two results in two specific areas. That is automatic translation of text, and automatic translation of images. So text translation can be performed without any preprocessing of the sequence, allowing the algorithm to learn the dependencies between words and their mapping to it completely new language.

(Refer Slide Time: 42:15)

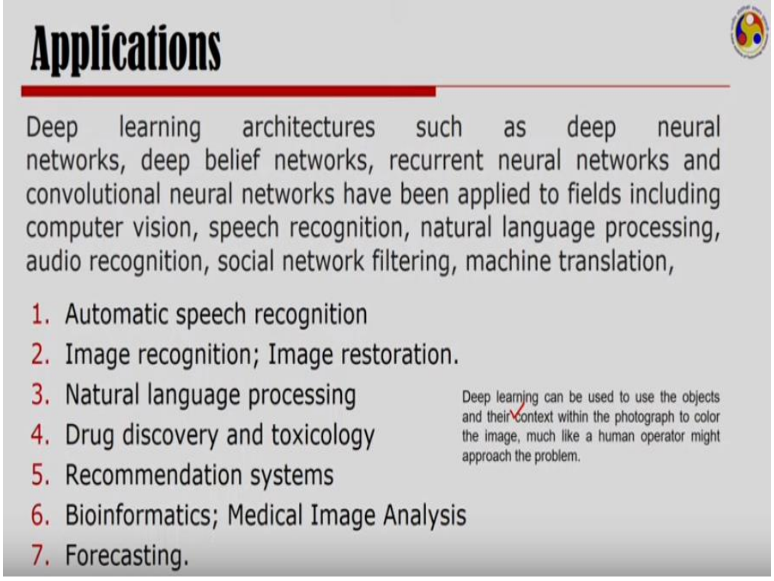

Deep learning have been used for object recognition, and have been used to learn the context of objects within photographs. Or to color the images, much like a human operator might approach this problem. And therefore, one of the most interesting applications of deep learning has been conversion of black and white photographs to colored photographs automatically, or even converting black and white movies to colored movies.

In fact, there are also applications where silent movies, audio has been included into the movie using deep learning architectures.

(Refer Slide Time: 43:11)

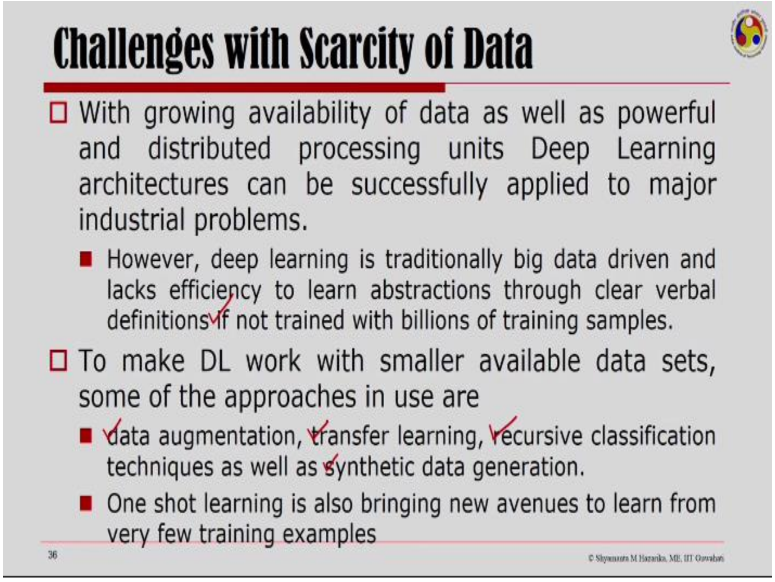

So finally, let us look at what are the challenges of deep learning. With growing availability of data, as well as powerful and distributed processing units, deep learning architectures are already successfully applied to many industrial problems. However, one thing to realize is that deep learning is traditionally big data driven, and lacks efficiency to learn abstractions through clear verbal definitions, if not trained with billions of training samples.

So if deep learning has to survive, I think it is pertinent to make deep learning work with smaller available data sets. And some of the approaches are already on its way in this direction, data augmentation, transfer learning, recursive classification technique, as well as synthetic data generation are looking at this problem. One shot learning is also bringing new avenues to learn from very few training examples.

(Refer Slide Time: 44:32)

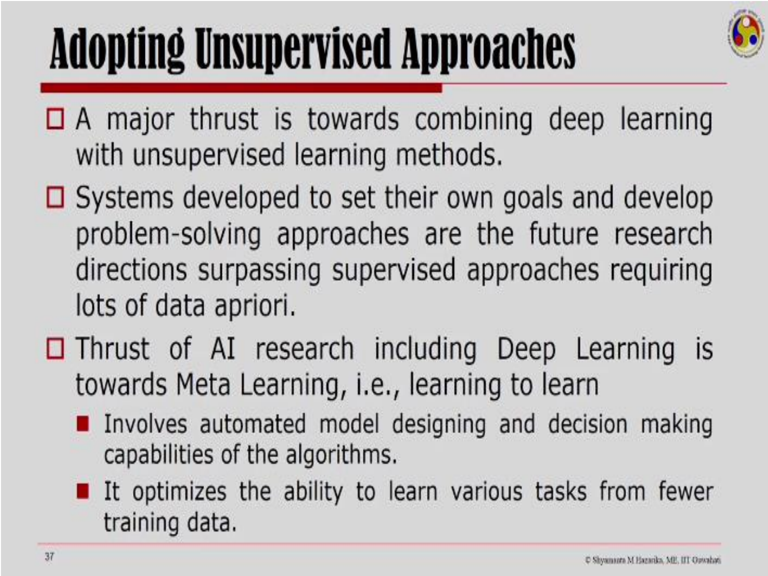

The other challenge is about adopting unsupervised approaches. A major thrust is towards combining deep learning with unsupervised learning methods. Systems developed to set their own goals and develop problem-solving approaches are the future research direction, surpassing supervise approaches, requiring lots of data apriori.

The thrust of AI research, including deep learning is towards Meta learning that is learning to learn. And this involves automated model designing and decision making capabilities of the algorithms. It optimizes the ability to learn various task from fewer training data.

(Refer Slide Time: 45:24)

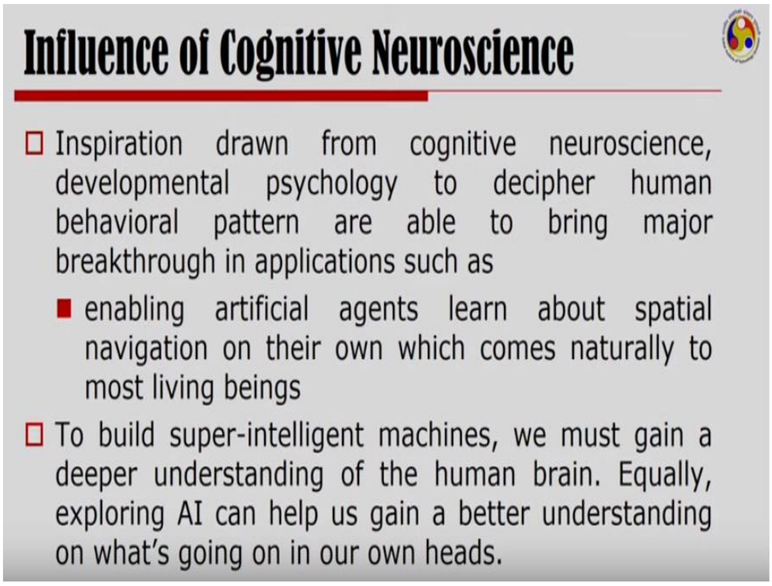

Finally, when talking of deep learning and deep learning architectures, one need to look at the influence of cognitive neuroscience. Inspiration that is drawn from cognitive neuroscience, developmental psychology, to decipher human behavioral pattern are able to bring major breakthrough applications, such as enabling artificial agents learn about spatial navigation on their own, which actually comes naturally to most living beings.

To build super-intelligent machines, we in fact, need to have a deeper understanding of the human brain. Equally, exploring artificial intelligence can possibly help us gain a better understanding of what is going on in our own brains. And therefore, this idea of cognitive neuroscience and artificial intelligence is intertwined, one drawing inspiration from the other. This is where we end this lecture on deep learning.

And this brings us to the end of module seven. And in fact, this is the last lecture of the course. So before signing off, let us look at what we have covered in this course, in terms of the learning objectives, vis-a-vis the different modules.

(Refer Slide Time: 47:18)

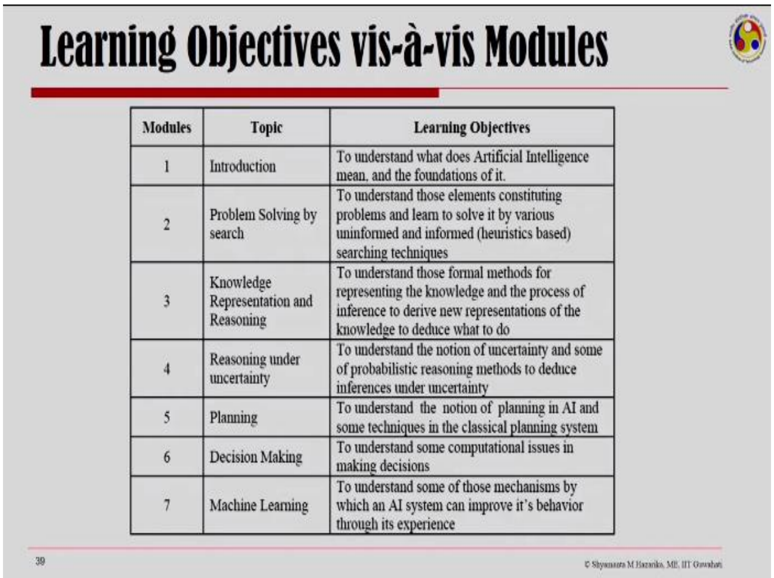

Learning Objectives vis-à-vis Modules

Modules	Topic	Learning Objectives
1	Introduction	To understand what does Artificial Intelligence mean, and the foundations of it.
2	Problem Solving by search	To understand those elements constituting problems and learn to solve it by various uninformed and informed (heuristics based) searching techniques
3	Knowledge Representation and Reasoning	To understand those formal methods for representing the knowledge and the process of inference to derive new representations of the knowledge to deduce what to do
4	Reasoning under uncertainty	To understand the notion of uncertainty and some of probabilistic reasoning methods to deduce inferences under uncertainty
5	Planning	To understand the notion of planning in AI and some techniques in the classical planning system
6	Decision Making	To understand some computational issues in making decisions
7	Machine Learning	To understand some of those mechanisms by which an AI system can improve it's behavior through its experience

So if you recall, we had our first module, which was on a broad introduction to the area, where we had traced the history of artificial intelligence. Looked at what are the dimensions of AI and then tried to get to the very crux of the problem in terms of what is strong versus weak AI. In the second module, which was on problem solving, by search, we looked at the artificial intelligence technique and then looked at production systems.

Therein we introduced both informed and uninformed search techniques. We looked at algorithms such as A star and AO star. That is we had looked at and or graphs and then we looked at constraint satisfaction problem and also how to search game trees. Our third module was on knowledge representation and reasoning, where after a brief review of propositional logic, we did cover first order logic.

We looked at the inference in first order logic, procedural reasoning and we looked at answer extraction in first order logic. Thereafter, we looked at reasoning under uncertainty in our fourth module and have looked at Bayesian networks and decision making under uncertainty using utility theories. We had our fifth module on planning, where we looked at partial order planning and planning graphs and graph plan.

In our sixth module on decision making, we looked at practical planning and two very interesting concepts on decision making. One, which was about sequential decision problems, how you make decisions in sequential decision problems. And we looked at the Markov decision problem. We thereafter looked at decision making, which involved making complex decisions.

We also looked at how partially observable Markov decision problems could be addressed using what are called decision networks. And then, we arrived at our final module on machine learning, where we provided a basic overview of the area of machine learning. Covered learning from observation including understanding how to learn a decision tree, get to know how to perform linear regression.

And then we looked at support vector machine, a supervised learning algorithm, possibly one of the most popular and one that needs to be in your toolbox, if you are talking of machine learning. And then, we introduced unsupervised learning, k-means clustering and hierarchical clustering. The final week had three lectures, one on reinforcement learning, which needs to be seen as a continuation of our discussion of sequential decision problems from module six.

And we had two lectures where we rather given exposure to the area of connectionist AI or in more plain terms, we introduced artificial neural networks and covered deep learning. Now this brings us to the end of this course. I hope that this course have been of some use to you. And you could take up AI in a big way with the fundamentals that we have covered in this course. Thank you very much.