Social Network Analysis
Prof. Shivani Kumar
Department of Computer Science and Engineering
Indraprastha Institute of Information Technology, Delhi

Chapter - 09
Lecture - 01

Hello everyone. Welcome back to Social Network Analysis. Today we will be learning a bit about deep learning. We will be diving into some of the basic architectures and some overview objectives of the deep learning; of the deep learning paradigm, and we will be going over some applications of it and how the learning actually happens in such kind of architectures. But, before we begin this deep learning introduction, we must know what deep learning actually is.

(Refer Slide Time: 00:57)

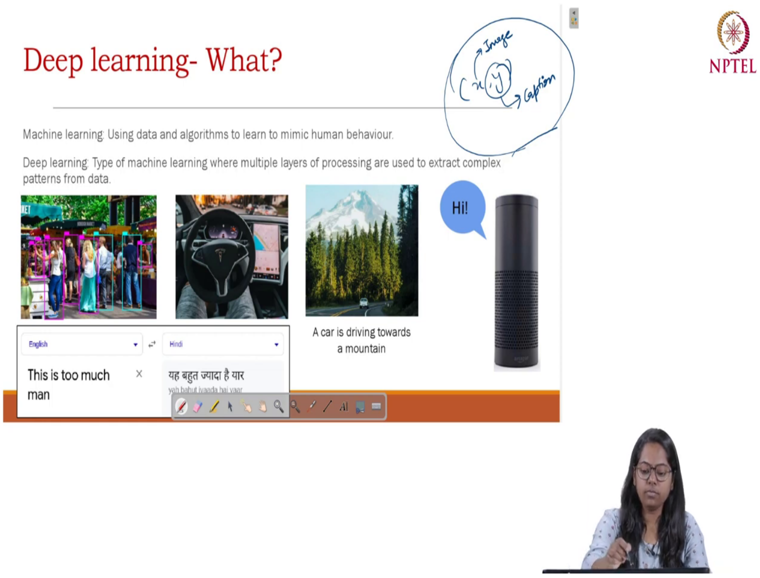

But again, before we know what deep learning is, we must know what AI or artificial intelligence is. Now, AI it is basically a field of computer science, where the researches that are involved in this field they want to make a system such that it can mimic human behaviour. That is it can make decisions just as a human would. Now, this whole task of AI, it can be used for various applications. For example, sentiment analysis or say translation problems.

Now, in order to achieve this bigger goal of AI, there are different strategies that can be applied to it. For instance, we have machine learning. Now, in machine learning basically there are some statistical methods or shallower networks, which are used to learn some kind of patterns in the data in. So, that we can map the given input to a desired output; now, as said before this desired output can be for instance sentiment classification or say emotion classification or dialogue generation or machine translation.

Now, in the case of machine learning, we have as said before we have some shallow networks. So, these networks they are they may capture some patterns in the data, but they are unable to capture highly complex patterns that may be present in the data; which even we humans cannot see if we just look randomly at the data. So, here comes the concept of deep learning.

So, basically in deep learning, we have multiple layers of processing which helps us to capture more complex patterns that simple statistical machine learning algorithms might not be able to do.

Now, this use of deep learning that is to achieve the main goal of AI. We see the use of deep learning almost every day in our surrounding. For instance, for the task of object classification deep learning is used. So, for instance we say we are given an image and we have to identify what all different objects are present in the image. So, for that purpose deep learning algorithms might be used. For instance, here we can see that given a given an image of say a marketplace.

We have to identify the people in it. So, here deep learning algorithms must be used to identify such instances. Next, we are aware that the they this is a era for automated driving right. So, in the case of automated driving as well deep learning algorithms are employed to learn how does a human drive. So, that that kind of pattern and that kind of behaviour can be mimicked can be replicated.

So, that a machine can learn how to drive on its own and then can just like be on the road and drive and avoid all the obstacles like a human would. Further, an interesting task that deep learning can help us achieve is basically image captioning. Now, in image captioning the input is simply an image and what we want to generate is a caption or a description for the image, which can help us to basically describe the image.

So, it identifies the different objects in the image it identifies how the different objects are interacting with each other and then it helps us to generate a caption which describes this image. Now, this caption though it might be useful for say a magazine or a newspaper kind of an industry, but it is also useful for say differently abled people. So, for instance, if we are if there is someone who is blind, who wants to understand what kind of image is being displayed on a screen on or on any app.

So, they can be they would want a caption or a verbal description of the image right. So, in that scenario also this image captioning can be used. So, another application for deep learning would be that basically we see almost every day would be the use of chatbots. So, we are so habitual in interacting with Alexa, Google or Siri in our day to day activities, that we oftentimes ignore how the background of these dialogue agent these chatbots work.

So, basically in the background deep learning algorithms are used. So, when we are; when we are querying these dialogue agents for instance Alexa, then we are firing up a deep learning algorithm inside of it, which basically identifies the intent the our intent, our objective that we want to that we want Alexa to achieve, that we want to achieve ourselves and then it queries back to its database and the internet that it has.

And then it comes up with a reply to us. Another day to day activity that involves the application of deep learning is translation. So, we often use this translation blindly or we often does not appreciate the amount of algorithms that are being used at the back end of it. So, in translation as well deep learning algorithms are used, where the input is in one language and the output is in some other language.

So, now, all of these; all of these instances of application that we just discussed like object classification, automated driving, image captioning, dialogue agents and machine translation. All of these the training part of these basically involves an input and a desired output pair. So, an x y pair right. Where x is the input which is like for the first case of image classification, it is the plain image without any classes without any objects being identified or in the case of say image captioning it is just the image.

And y is the desired output for like in image captioning, it is the caption. So, if the in the case of image captioning, this x is the image and y is the caption. Now, for all of these 5 cases, we have some x and some desired y. So, where we are aware of the y that we want like that is the label or the ultimate goal that we want, this kind of learning and like we provide this x y label to our learning algorithms. Then this kind of learning, it is known as supervised learning, because this learning is being supervised by this ground truth labels of y that is what we want.

(Refer Slide Time: 08:23)

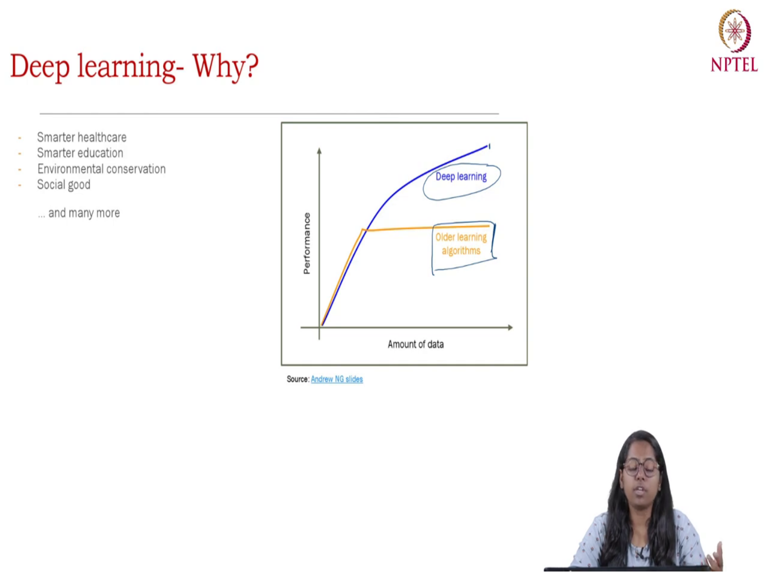

Now, we know what is deep learning, but why is it necessary to understand and learn deep learning. So, we already saw quite a few applications of deep learning, where it is used and there are hundreds of more applications where deep learning has become the daily norm. But, again the one more important thing that we must know is the that the existing older algorithms of say machine learning or other statistical methods.

They often they perform good, but often at a particular amount of like after a particular amount of data that is being given to these algorithms, the performance of these algorithms that often plateaus that is even if we increase the data after a certain amount, certain point the performance does not increase much right. So, basically we can say that these algorithms they are unable to exploit the amount of data that we have nowadays.

So, nowadays there is a abundance of data a plethora of data is available, but these older algorithms they are just unable to exploit it. Whereas, with the availability of more data, we can use these deep learning methods which can actually exploit this advantage of more data to give us better performance of for any of xyz tasks that we want. Now, this performance increase because of the amount of data that we have.

So, we know that these deep learning algorithms specially, these supervised deep learning algorithms they are data hungry algorithms that is more the amount of data that we give to these algorithms or these methods the better performance we will get. Now, this deep learning methodology it can be used into like for example, in their day to day activities it can be used into various domains; for example, in healthcare.

So, deep learning in healthcare is a major thing that is booming nowadays. For instance, during the current pandemic of Covid 19, we saw that using deep learning methods doctors and researchers they were actually able to identify the intensity of the Covid infection into the patients by just looking at the X-ray of the patient. And if we pass that X-ray image to say a deep learning algorithm then that algorithm was able to identify the intensity or the chances of say hospitalization or fatality in a person.

Now, definitely other than the pandemic other than Covid 19, the healthcare domain is using deep learning in a in a large manner; for instance, say a cancer detection at an early stage by looking at the MRI or the PET scans of the patient and other such stuff. Now, apart from healthcare, education is also a domain where deep learning is affecting a lot.

For instance, we might have say like smart teachers or smart examinations, where the students answer sheet is checked automatically or we can have smart plagiarism detection where like the plagiarism between 2 students or the student and the existing literature on the internet is being checked automatically using these deep learning algorithms.

Now, apart from these deep learning can also be used for other socially relevant topics. For example, say environment conservation or other social good for example, identifying fake news or identifying hate speech and even combating such activities. Now, and definitely apart from just these few application that we mentioned, there are a lot more applications where deep learning is actually used in daily way.

(Refer Slide Time: 12:38)

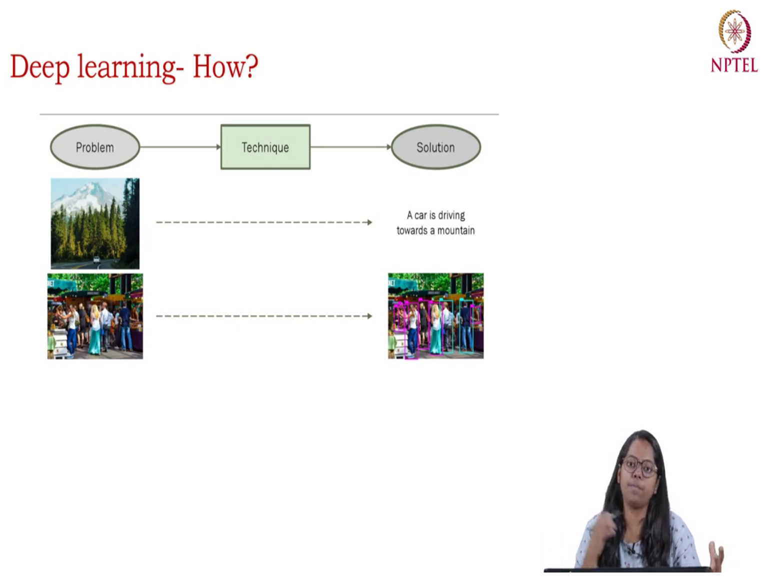

Now, we saw what exactly is deep learning. We saw why do we need deep learning, but now the question comes, how can we actually apply deep learning to such methods? So, all of these problems, all of these scenarios that we just saw, they basically have some kind of input; we which we might call the problem and an associated solution to it, so for example, in the case of say image captioning.

The problem or the input is the image and the desired solution is basically the caption of the image. And in between thus there is some kind of algorithm, some kind of technique that is being employed which basically helps us to convert this problem of say the input the image to the solution that is the caption in this case. Now, another example for the case of image classification, the problem or the input is the simple image without any without any class objects and the solution is the bounded boxes of the different classes or the different objects that are present in this image.

(Refer Slide Time: 13:54)

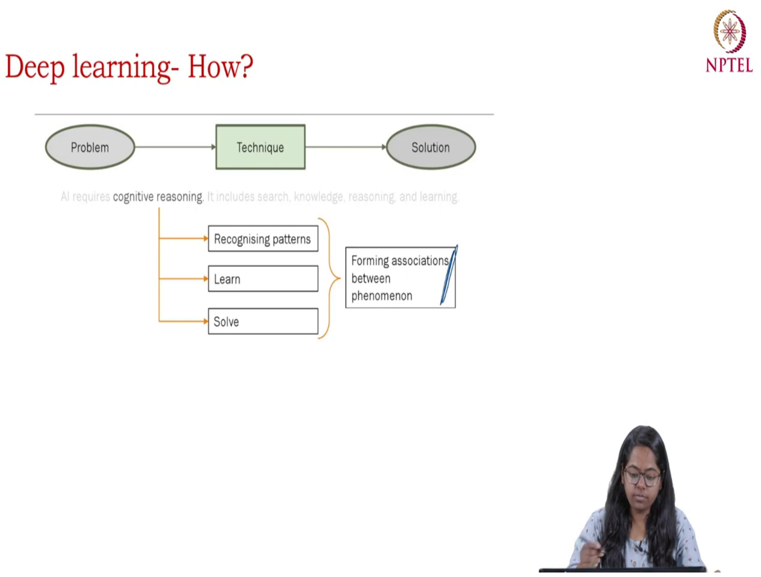

Now, how to do this? So, now basically, deep learning or AI it involves a some cognitive reasoning that is happening in this technique part between the problem and the solution. Now, this cognitive reasoning it basically includes searching that is based on all the previous, all the past experiences that we have we search what kind of experience the like for the current experience resembles the most too.

So, for example, like we human beings we live our lives and we encounter various different situations various different various different experiences in our life. And whenever we are encountered with a new problem new experience, we know what should be our next step based on our previous experiences and previous results of those experiences and the other steps that we took.

Similarly, we want as AI is all about mimicking human behaviour. So, on the similar aspect in this in the computer science domain also, if we want AI we must be able to search our previous experiences or our previous like every historical scenarios that we have seen. So, that we can have some kind of knowledge about the current situation and the different pattern that we are observing right now.

And we can reason based upon this knowledge and then we can take an action and moreover, based on the outcome of this current action and the scenario that we were encountered with we are learning for the future. That is if we are if we encounter such a scenario again in the future, we know that what the consequence of the action that we took at that time would be.

So, basically cognitive reasoning or this technique part between the problem and the solution involves searching, then obtaining knowledge, then reasoning based on then knowledge. and then learning based on the outcome of our actions. Now, this cognitive reasoning it basically involves recognising patterns. That is so for instance, if we compare it with human beings; we can say that a pattern is something that like a scenario that we observe on a daily basis for example, if we see a baby crying we can associate this pattern with say the baby is hungry.

So, we can associate a crying baby to hunger. So, this is basically associating a pattern to a like a reason or a knowledge that we have. So, we want this similar kind of pattern associations happening in the deep learning system as well. So, the first part of cognitive reasoning is to recognise pattern, then after recognising this pattern then learning what to do in such a scenario. For example, if we see a baby crying, we know the pattern resembles to that the baby is hungry. So, we give the baby some food.

Then finally, solving that is what we just mentioned that we know based on our past learning we know that the baby is hungry. So, the solving the solution would be to give the baby some food. So, the cognitive reasoning it basically involves 3 things, recognising patterns, then learning based on these patterns and then providing us some actions, which helps us solve the problem being identified.

So, basically all of these 3 things, it constitutes of forming associations between say the input or the input signals that we have the scenario that we see and the corresponding actions that we should take. So, that is it is basically help us forming associations between different phenomenas.

(Refer Slide Time: 17:51)

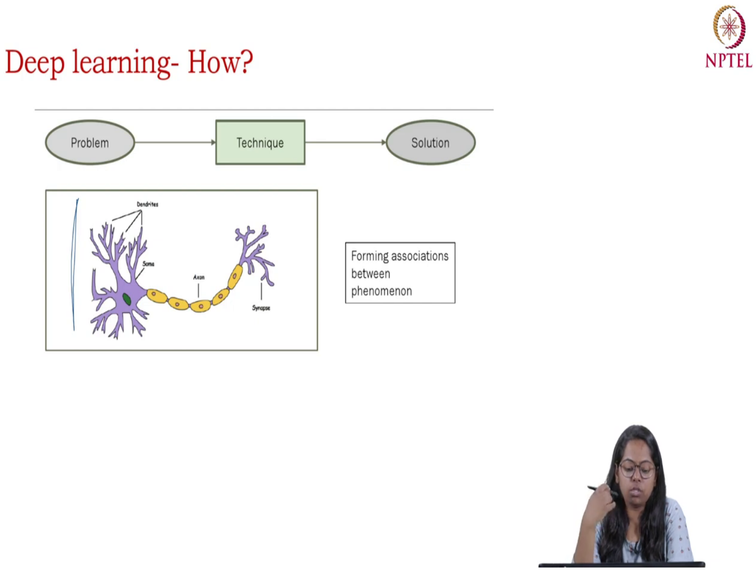

Now, in the case of human beings these associations these are being made by in our brains right. So, but what exactly constitutes the brain? So, we have these kinds of neurons in our brain. These neurons or brain cells that basically are they are called, we have a series of these neurons a lot of these neurons that are present in our brain which actually helps us to perform this cognitive reasoning that we want. That is it helps us to understand the scenario around us, it helps us to understand what kind of past experiences this current scenario resembles the most.

And then it helps us to take a decision based on all the understanding of the environment that we obtain. Now, these multiple neurons that are present in a brain, they are fired based on some stimuli that we receive right. So, for instance, in the case of again the baby crying we see a baby crying by our eyes right and we hear the crying by our ears.

So, the these 2 stimuli helps us to activate some neurons in our brain, which then understands that the sound that we are hearing it is actually of a baby crying. So, but obviously, in a computer system these kinds of neurons they are not present.

(Refer Slide Time: 19:20)

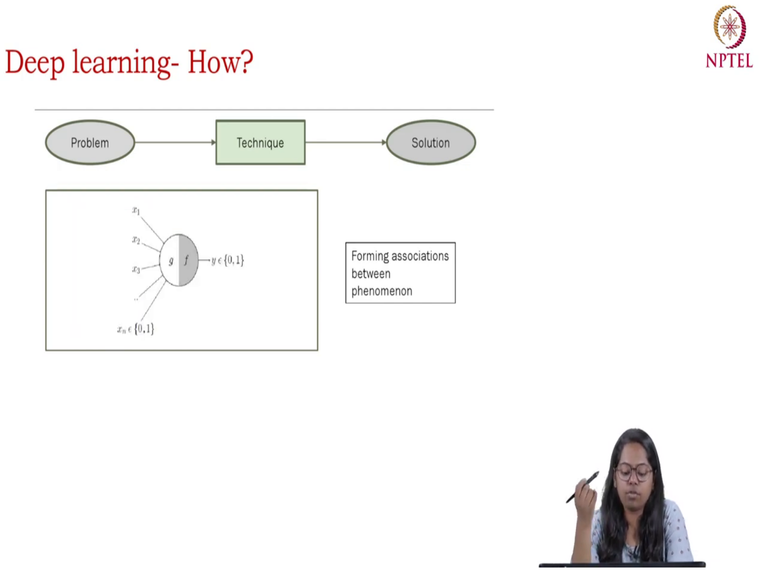

So, various different researchers; they have tried to give some mathematical models of these neurons. For example, so, we will go into these models like we will go into a few of these models, but the basic functionality of these mathematical models, it is to obtain a similar goal that a neuron in our brain or the brain cells actually does. That is to form associations between different phenomenas.

(Refer Slide Time: 19:51)

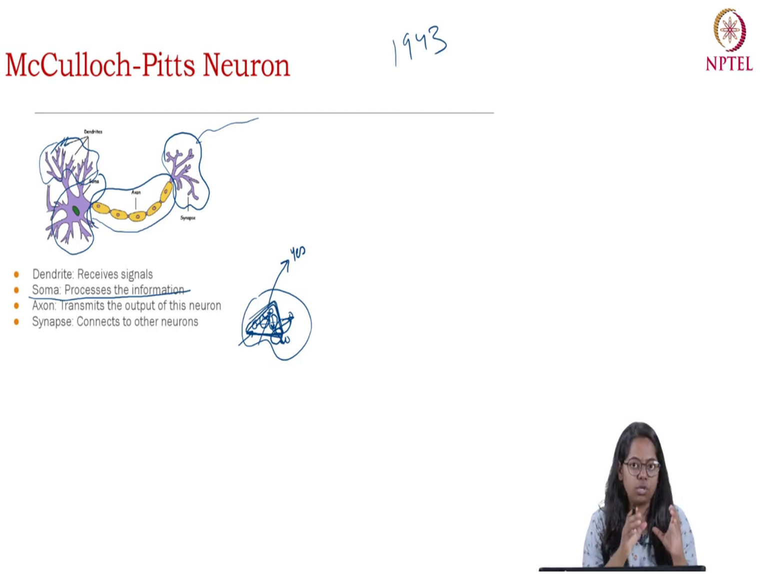

So, we will study the first neuron that was proposed in 1943 by McCulloch and Pitts. So, in the normal biological neuron, what so, there are basically 4 most important aspects of this neuron or 4 most important parts of the neuron. For example, the first part is the dendrite.

Now, these are the dendrites where basically these are the ends of the neuron, which receives the signal from say the other neurons or from some stimuli.

For example, the eyes or the or a our nose or our ears. Then we have the soma, that is the body of this neuron, which basically processes this information that is received by the dendrites and comes up with; and comes up with the decision, whether this particular neuron should be fired or not that is whether this neuron. So, whatever the function of a particular neuron might be, for instance in our brain suppose, this is our brain; we might have different neuron for different purposes right.

So, we might have these 3 neurons, which basically decide whether the sound that we are hearing right now is of a baby crying or not. So, if we say the pitch of the sound is high; we say that yes the baby might be crying. If we say here that the like the there are there is some kind of tears in the baby's eyes, then we can say yes the baby is crying. And based on these observation these 3 neurons the aggregated observations of these 3 neuron the processing of these three neurons we can come up with a with the final objective that yes the baby is crying.

So, this processing that is happening inside each of these neurons. This is happening in this in the neuron body that is called the soma. Then we have the axon which is just the transmission part of the neuron, which transmits whatever output this soma the neuron body has generated to the other end of the neuron, which is called the synopsis which basically connects to the to the further neurons in the system, in our brain.

So, for example, these here these all neurons are connected to each other. So, this connection they are basically happening via the synapses and the dendrites. So, these synapses they would be connected to the dendrites of the further neurons right.

(Refer Slide Time: 22:52)

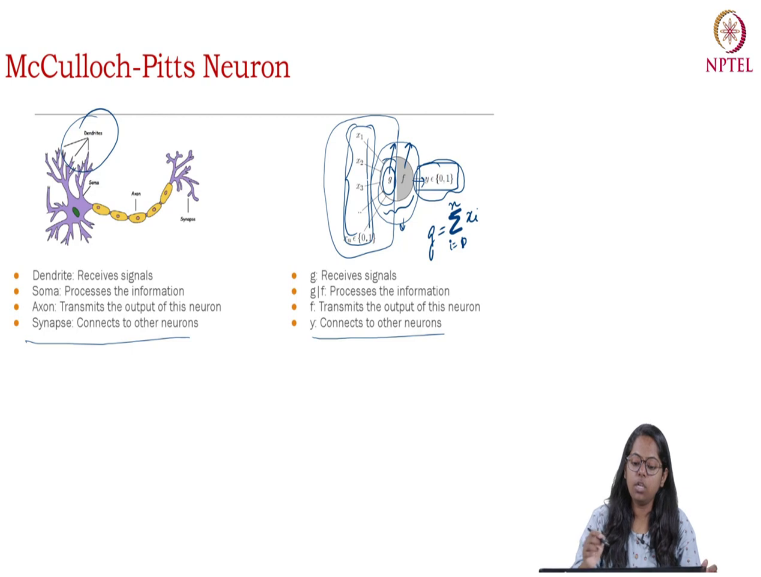

Now, congruent to these 4 aspects of the biological neuron McCulloch Pitts they actually came up with a mathematical model of a neuron mimicking these 4 attributes in their model. So, in this model we have some inputs, these xs, which are basically binary inputs that is each of the input can take up a value between either 0 or 1. Then these inputs they are passed on to this processing unit, which have some g and f. So, these are basically 2 functions, which here g is the aggregated sum of these inputs.

So, here g is basically the sum of all the x is, where i belongs to say one here till n and f is then the decision making the decision making aspect of this whole processing unit. And then finally, this decision is passed on to the like to the further neurons or the final the decision is considered the final output. So, if we are to compare this mathematical model with the biological neuron, then we can say that these inputs these x 1, x 2, till x n they are basically like the dendrites right.

So, these are these are the inputs that receives the signal to this g part. So, this whole part can be considered as the dendrite whereas, this g and f module this basically this processes this information the input information and decides whether the neuron should be fired or not. So, this acts as the soma part of the biological neuron. Whereas, the transmission happens here that is the axon and then finally, we have the output which can be passed further to the next neuron or considered as the output, which resembles the synapse from the biological neuron.

(Refer Slide Time: 25:01)

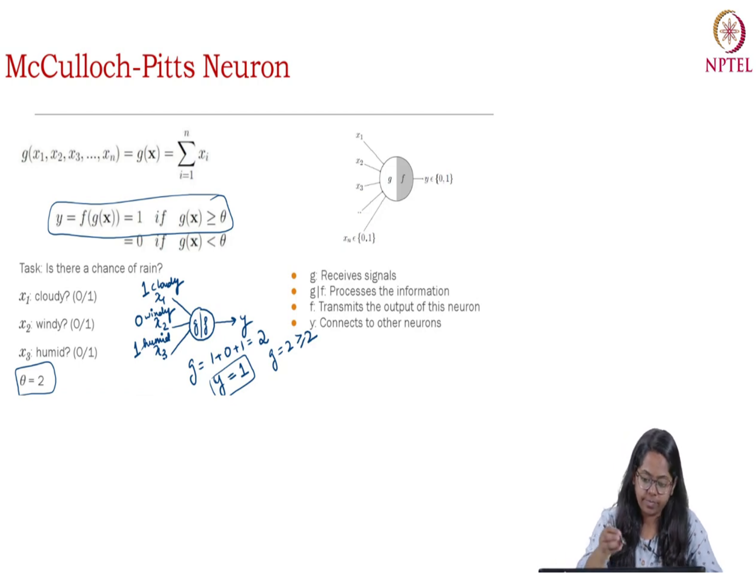

So, now, this McCulloch Pitts neuron model, it as mentioned before it receives signal it processes this information; then it transmits the output of this neuron and it further connects to other neurons or the final output. Now, how do we take a decision of whether this neuron should be fired or not based on the inputs. So, this is done by something known as the thresholding function.

Now, in the thresholding function, what we do is that we find out the aggregated sum of this binary inputs that we have. And based on the total sum of these inputs we decide whether the output of this neuron should be 0 or 1. So, again the inputs and the outputs of this McCulloch Pitts neuron is basically binary. Now, here in this equation it just it is just a mathematical modeling of what we just mentioned.

That we have these inputs xi, which are all binary and we perform a summation over all of these inputs and this is this summation this whole aggregated sum is called as g x that we have here. Now, we pass this g x this whole sum to this function f, which is basically y this function which gives us a value 1, if the value of g x or the sum of all the inputs is greater than or equal to a threshold that we said.

If the value the sum is greater than this threshold, then we say that the neuron is fired or the value of y is basically 1, but if this value is less than this threshold, then this neuron is not fired and the value is basically 0. So, now we will understand this by an example for instance, say this we have one neuron with a task to identify that whether it will rain today or not.

So, in order to identify whether it will rain or not, we might have some attributes that we want to consider for example, we might have we might want to see whether the environment whether the weather today is windy or not, it is humid or not, or whether it is the sky is cloudy or not. And based on these the values of these 3 attributes, we might want to take a decision of whether it will rain or not.

So, here what we do is since these inputs they are binary that is whether it will be cloudy or not windy or not humid or not. So, for here we will have; we will have network like this. We will have x 1 x 2 x 3 the input which will go to our neuron which will be g followed by f and then we will have the output y. Now, suppose the theta that the threshold that we said we want to say that if any two of these 3 conditions are satisfied that is if it is cloudy and humid.

We say that it would rain or if it is windy and humid we say that would rain. Or it if cloudy or windy we say that would rain. So, that is if any of the 2 conditions of these 3 conditions are satisfied, then we say that we want to fire this neuron or we want to say that the output of this neuron should be positive should be a yes or a 1. So, that is why we set this threshold of theta 2 2. So, here suppose this x 1, it represents the cloudy attribute. x 2 represents the windy attribute and x 3 represents the humid attribute.

Then if the input comes such they that the environment is cloudy not windy and humid and we calculate this g. So, g is just the summation of these x i. So, it will be 1 plus 0 plus 1, which will be equal to 2 and according to our theta that is 2 and based on this formula that we have here y is it should be 1 as g x that is g that is 2 is basically greater than equal to 2 or equal to 2 right. So, that is why y here should be 1 that is in this particular scenario, we will say yes, it will rain.

(Refer Slide Time: 29:49)

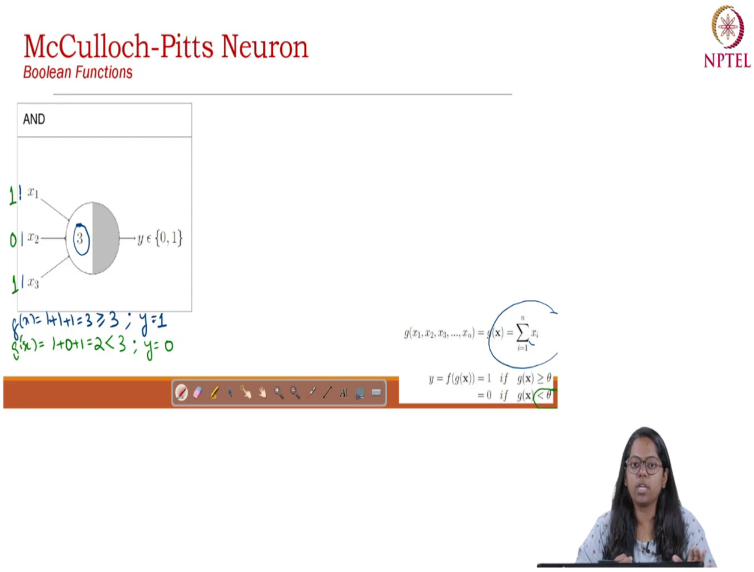

Now, since we saw that the inputs and the outputs are basically Boolean functions are basically 0 and 1. So, we can use these neurons to actually mimic the Boolean functions. And as we know if we are able to mimic Boolean functions we can use some like we can use the Boolean functions in a different manner to basically answer any question right. So, we look at some of the Boolean function that we have.

So, this part here, it represents the equations of our McCulloch Pitts neuron that is the summation of the xi, if it is greater than a particular theta particular threshold, then the neuron is fired and if it is not then it is not fired right.

Then we will just see the AND gate first; that is if all the 3 inputs are positive are 1; then only this gate will then the output of this gate will be 1 right. So, when the summation of this x 1 x 2 and x 3 that is equal to or greater than 3, then we will say that the y should be 1 correct. So, that is why here our theta is actually 3, because if x 1 x 2 and x 3 are both are all the three are 1.

Then only this neuron should be fired. So, based on the equation here if we look, if all of these three are 1, then our g x become 1 plus 1 plus 1 plus 1 equal to 3, which is basically greater than equal to 3 and that is why our y here would be 1. But now, what will happen if this is not the case. For example, if we have 1 x 1 as 1 x 2 as 0 and x 3 as 1, then our g x in this case would be equal to 1 plus 0 plus 1, which is equal to 2 which is basically satisfying this condition, because it is less than 3. So, here y will be 0.

And similarly we can do for all the values of the truth table and we see that if the threshold is set to 3, then we are able to obtain the correct values that we want for the AND gate.

(Refer Slide Time: 32:20)

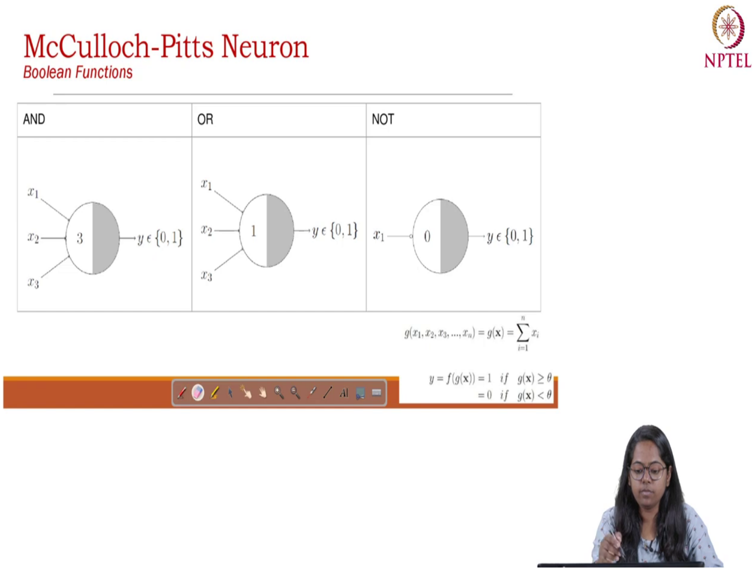

In a similar approach OR gate and the NOT gate can also be implemented. So, here in a very similar approach just like we saw the AND gate OR and NOT will also be implemented. So, right now we are not going into the details of it. So, the main objective here is to let you know that since, different Boolean functions can be implemented by using this McCulloch Pitts neuron. So, complex mathematics or complex decisions can actually be taken, if we are to put different neurons together to identify these kinds of functions.

(Refer Slide Time: 33:02)

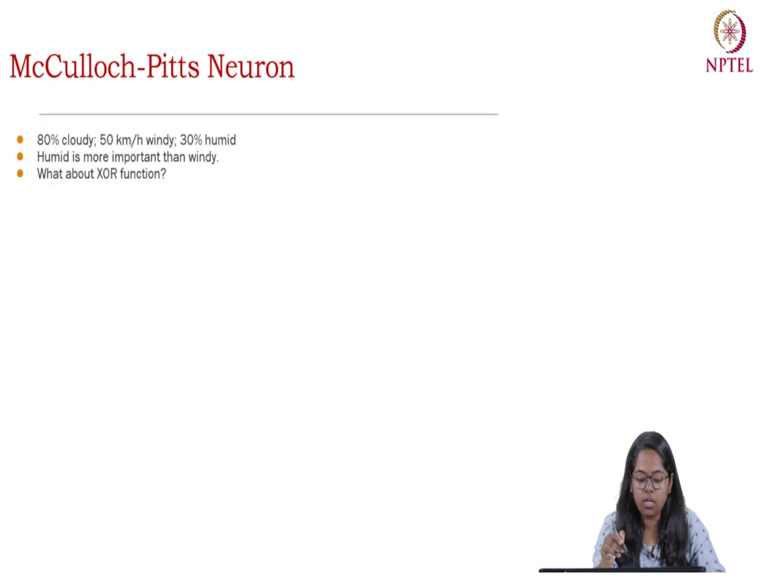

But there is some drawbacks obviously. So now, what will happen if instead of binary inputs, we have some continuous values as the input? For example, instead of just identifying whether it is cloudy or not on a single day, what will happen if we are to say that it is 80 percent cloudy today or the speed of the wind is 50 kilometers per hour.

Or the humidity the level of humidity in the environment is 30 percent today. That is it is not a binary input rather it is a; it is a continuous or a real valued input. Another important thing is that what will happen if one attribute of the input is more important than the other attribute. For example, if humidity participates more towards the chance of raining than the say than the wind and the cloudiness of the environment.

Then it will not be considered in the McCulloch Pitts model. Since, in this model all the input are weighted equally and are considered as equal whereas, in the real world scenario that might not be the case always. Another drawback here or another limitation of this neuron of this mathematical model is that the boundary the decision boundary that it provides is only for linear function that we saw that AND, OR and NOT they are basically handled by this neuron. But, what about the non-linear boundaries like the XOR function.

(Refer Slide Time: 34:53)

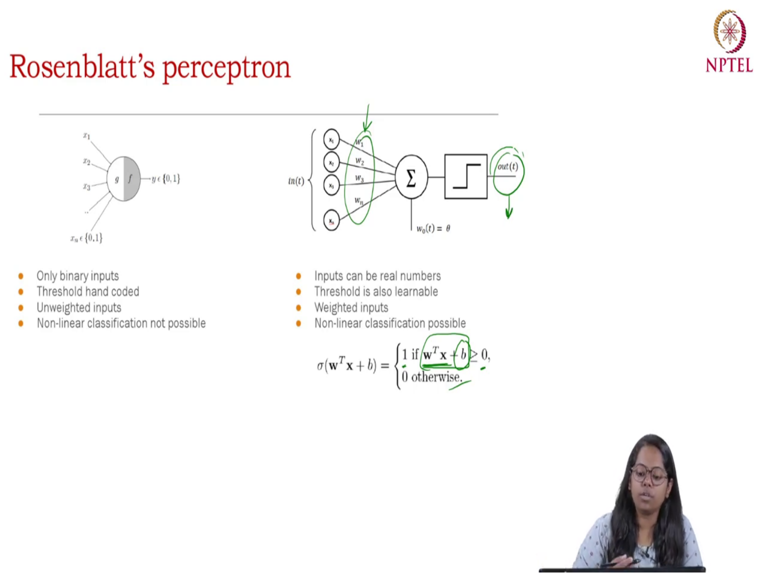

So, for all of these drawbacks, and more we will be reading about the next neuron model, that is basically the Rosenblatt's perceptron model, which is the; which is the first such model which helps us to which basically gives rise to the modern deep learning architectures that we have. So, in the traditional McCulloch Pitts neuron we had only binary inputs.

We had the threshold was also hand coded that is as we saw in the AND gate we decided ourselves so the threshold should be 3. So, this is like a hybrid parameter that we provide ourselves. Then the inputs they are all considered as equal, that is there are unweighted inputs and the non-linear classification it is not possible in the McCulloch Pitts neuron.

Whereas, if we look at Rosenblatt's perceptron, here we see that the inputs can actually be real numbers. Here the threshold is also learnable; we will see how this happens. Then here these all of these inputs which might be real numbers, they are associated with some kind of weights. So, the inputs are basically weighted and here non-linear classification is also possible, we will get into it now.

So, and this is basically the formula that is used for Rosenblatt's perceptron that is a weighted sum of all the inputs plus some bias that is the threshold that is to be learnt. If that is greater than or equal to 0, then we say that we are to fire this perceptron otherwise, it remains at the 0th state.

(Refer Slide Time: 36:34)

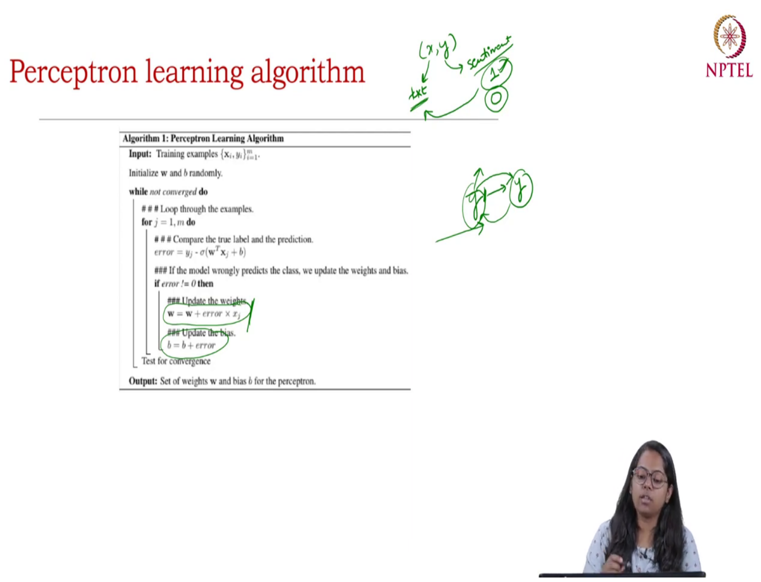

Then we have this perceptron learning algorithm. So now, this perceptron we know that it takes the real values as input. It has some kind of activation function which basically converts this weighted sum of the inputs into a decision that we want either 0 or 1. And then it helps us to take that decision, but now these values for these weights and this bias is they are to be learned right. So, they are maybe they are randomly initialized, but we want to learn these values.

So, that this output whatever output we are getting at every time, it is closer to the desired output that we want. So, for example, here we are like in this particular scenario, we are given some kind of x y pairs as the input. So, let us take the scenario for example, of sentiment classification. So, here x would be the text of which the sentiment is to be classified and y would be the required the sentiment, the sentiment class of this text x.

Now, it can happen that when we initialize the weights of the perceptron randomly, then for some sentiments that are actually positive in our in our grounds test they are being classified as negative right. So, now, we want to modify the weights of the perceptron in such a way that this classification, it is happening it should happen correctly for most of the part in our most of the instances in our input. So, in order to do that, we modify the weights and bias of our perceptron in such a way that the output moves closer to the desired output.

So, now we will go in depth of this perceptron learning algorithm, when we learn about back propagation, because here the weight update that is the way it is being passed being passed and the weights the this error, basically the error that is being used to modify the weights and the bias of the perceptron that is all part of the back propagation. Because in the forward pass we just calculate the value of y that is being predicted.

But in the backward path, based on this predicted y and the ground truth y that we actually have, we modify the weights and bias in such a way such that its predicted y moves closer to the ground truth y. So, we learn more about this algorithm in the back propagation part. Next we will so, today we saw the perceptron, we saw the basic neuron and we also saw why and how do we perform deep learning.

(Refer Slide Time: 39:55)

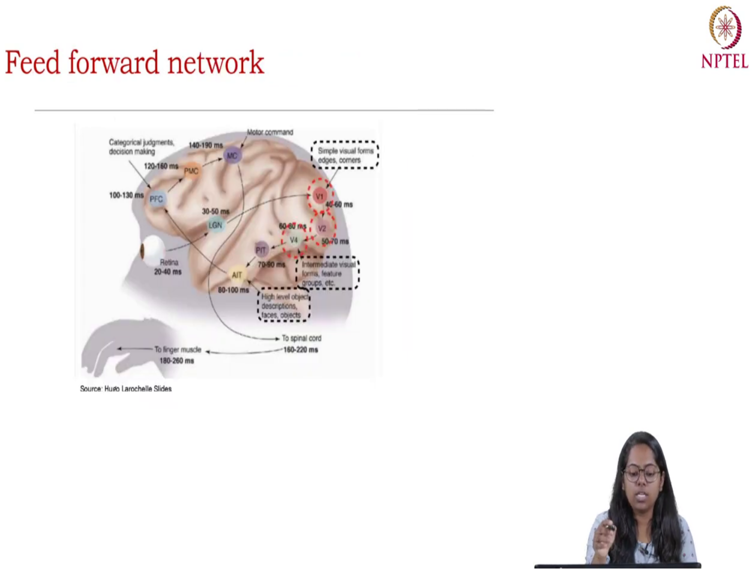

In the next lecture, we will see the feed forward network and the back propagation algorithm in order to understand how the learning actually happens in a more complex manner.

Thank you.