Machine Learning for Engineering and Science Applications
Professor Dr. Ganapathy Krishnamurthi
Department of Engineering Design
Indian Institute of Technology, Madras
Introduction to Generative Model

(Refer Slide Time: 0:15)

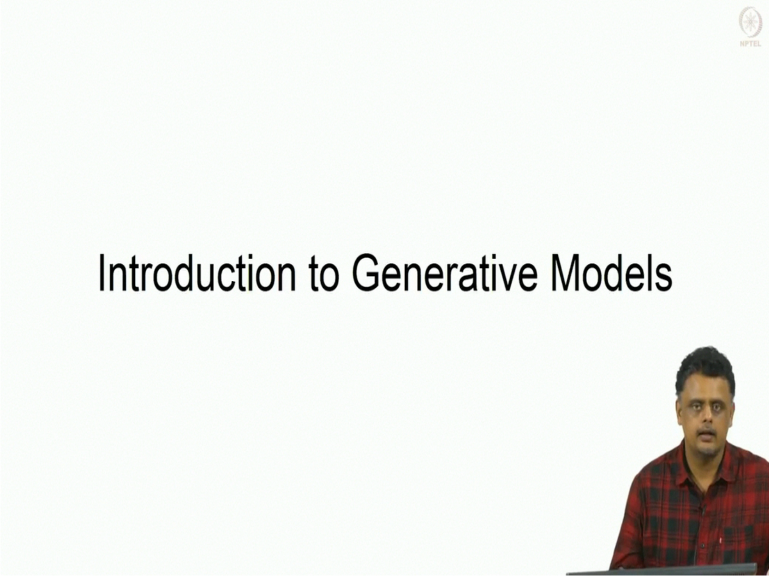

Hello and welcome back, in the next series of lectures we look at generative models, especially the deep generative models and in this video we will just take a look at what they mean and what they are used for.

(Refer Slide Time: 0:28)

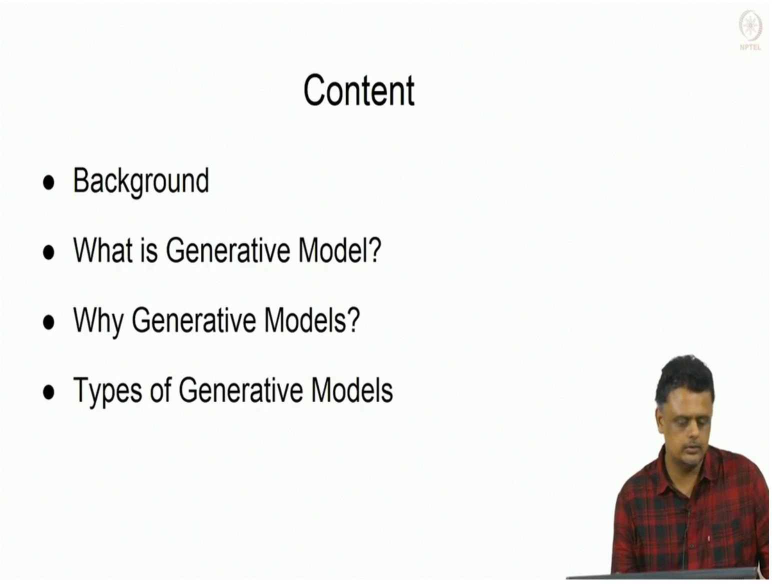

So is a brief outline of this lecture, so we have some background to cover and see what are the generative model is and why we need them and the types of generative models will be covering as part of this course okay.

(Refer Slide Time: 0:42)

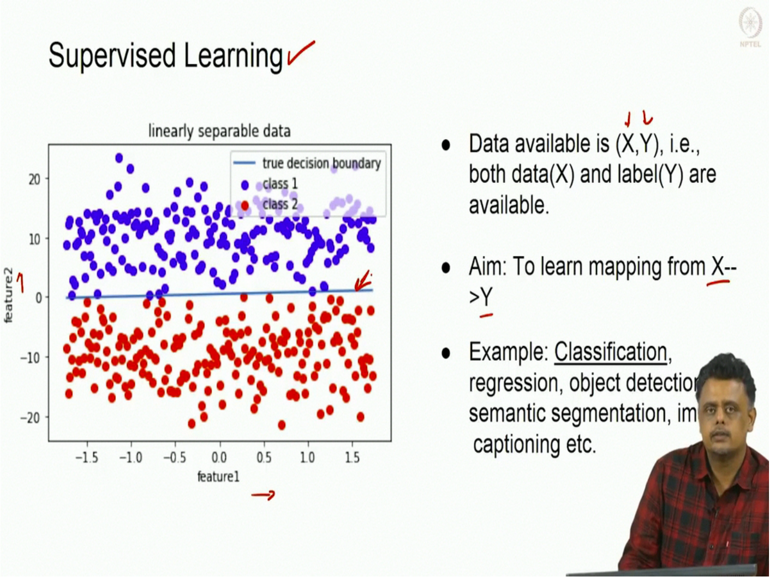

So, so far, you have looked at a problems wherein we have this spares available X and Y, where access the data, it will be the multidimensional data, images, whatever form of data that you are having, that you have with you and the label associated with it, so typical task is a classification task is yes or no or like we saw a imaginate data you have 1000 classes, so you have a input X and the corresponding label is given, so these are label their assets and what we use the deep neural networks for is to learn the mapping from X to Y okay.

So X is the input and if we have a deep neural network which processes the input and output say probability score which we threshold and make that into a class label right, so this is in the context of supervised learning right, this is what you have seen, so another way of looking at it is, that the deep neural network defines a classification boundary like shown here in this graph, wherein we have two features, future one and feature two this could be some arbitrary numbers have to scaling.

The red and the blue dot are the two, or the datasets or the data points corresponding to the two different classes and what the neural network does is to determine this boundary if you can look at it that way and for which we have access to both the data, as well as the label.

(Refer Slide Time: 2:18)

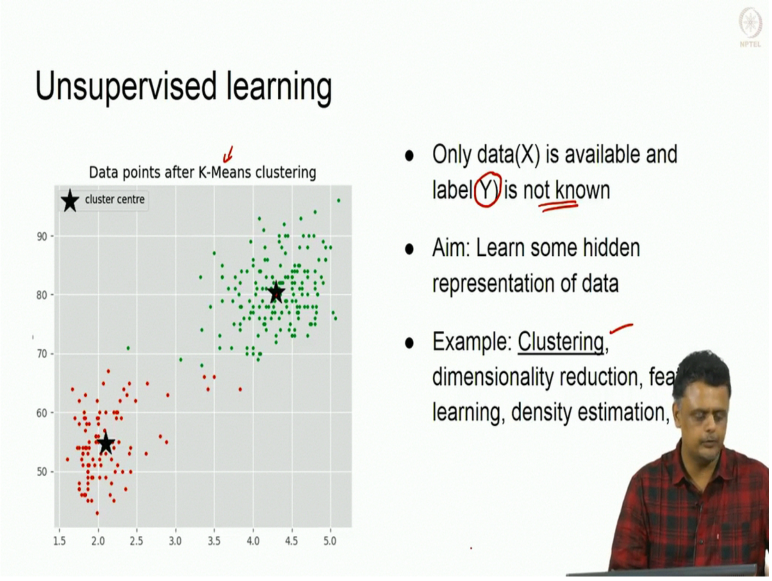

We have also seen that, you know unsupervised learning wherein there is no associated label, so Y is not known right, Y is not given and what we try to do here in this situation is to learn some underlying structure in the data right, so that falls under unsupervised learning, some of the techniques we can also call them clustering technique is what typically use in that context, so, for instance, K means clustering wherein, which is specify the number of clusters that you hope to find in the data and you end up with a labelling accordingly right.

So in this case we have again two features and what came in algorithms does as is to, if you think that there are two cluster present in this data, the K means algorithm determines the centre of those clusters right and is able to label each of the data points based on the distance from the a cluster centres, as belonging to either one class or the other. Okay, so here we are just trying to determine some underlying structure to the data right and we do not have access to the labels of the data, we just have the raw data available and we try to find some underlying structure alright, so, in this context, classifying is what we have seen, feature learning or density estimation all these problems fall under unsupervised learning.

(Refer Slide Time: 3:52)

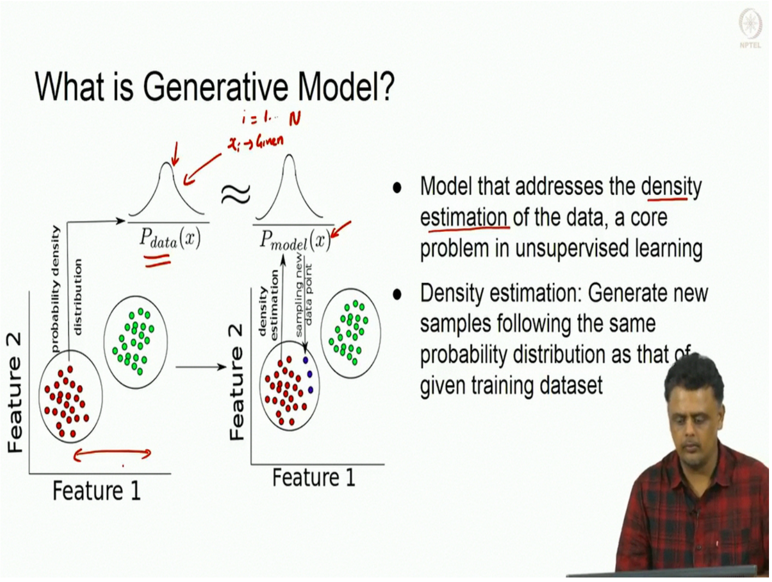

So part of this unsupervised learning is the density estimation problem. Okay, so it falls under unsupervised learning and the idea is to determine probability density model for a data from which you can draw new samples okay and you determine that probability density, your model density from the given data density, so you have given data X, so this is, let say that you have N data points, 1 to N data points and so that is basically XI are given, which is an approximation to your, from which you can get an empirical density right, that is P data of X okay.

So we do not know the true P data of X for which you can infinite number of points, so for all possible access that you get your hands on or that exist which is not possible, so you with the data available, you can given empirical density estimate of the given data, but we what we want is to determine a model from which we can sample the data points okay, so which is basically the generating sample, so hence word generated, terminology generate to models right.

So given probability, so given this data points okay, so again we would not have access to the labels for say some most of the time we just have the raw data and we have the index data and what we want to do is to figure out or have model for the underlying probability distribution which is that is what you want, so that we can draw new samples from it, which are similar to our training data. Okay, so we assume that our training data has representative of the problem that we are trying to solve okay.

(Refer Slide Time: 5:57)

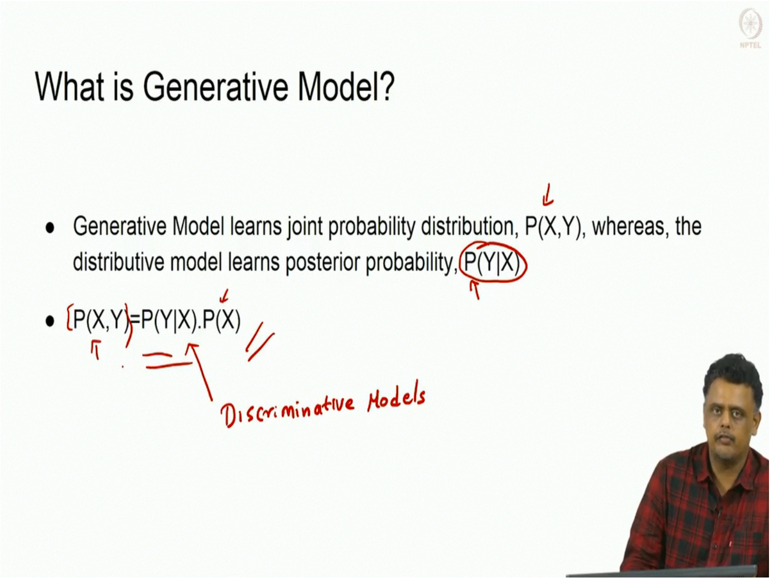

So why this useful and do we do with it, so before we proceed there we can also, there is another way to look at it, which is to say that the generative models learn the join probability distribution P of X, Y, whereas the your classifier for instance right, your if you are using a deep neural network to classify data into one of K classes and the output of neural network, the scores that are output by neural network can be interpreted as the posterior probability of the class given in the data, so that is P of Y given X okay.

So that is a, so that is not a generative model right, so what we actually want is this, of course this are related through base rule here right, so if you have a good model for a prior P of X and we can actually figure out this, so this is one way of looking at it, so what we are, so this is what is called a discriminative models right, this are called discriminative models, in the context of classifications, so what is the probability of label even in the data that is what are typical machine learning algorithm outputs.

So to most of your deep neural networks that you use for classification but we do seek something of this sort where we have, where we actually want to model the underlying data distribution that is P of X, typically is what we want to get.

(Refer Slide Time: 7:38)

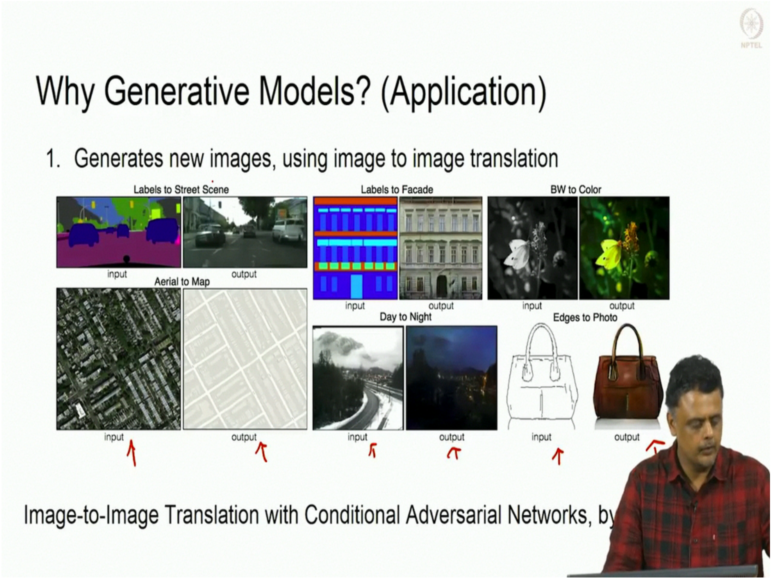

So why do we need relative models? Okay, so let say basically we are now trying to look at what kind of applications can they possibly serve, so the most common application is they can generate new nice pictures right, we can also use them for image to image translation, there is shown in this examples or for instance, we gave them labels or outlines can you generate new data, so, for instance, blacken into transform black-and-white to color

So it gives the aerial map of a city, can give you this, you know this kind of output as you see in Google maps or difference between night and day, here day and it has the same night time and if you give labels, once again, if you train the models, so that if you give it an outline you can give a output that you see it is a new design for handbag if you can think of it that way, so there is aerial to map, you know labels to street scenes, labels of a facades, the black-and-white to color, so these are some need cool applications that you can think of if you say a generative models alright.

(Refer Slide Time: 8:50)

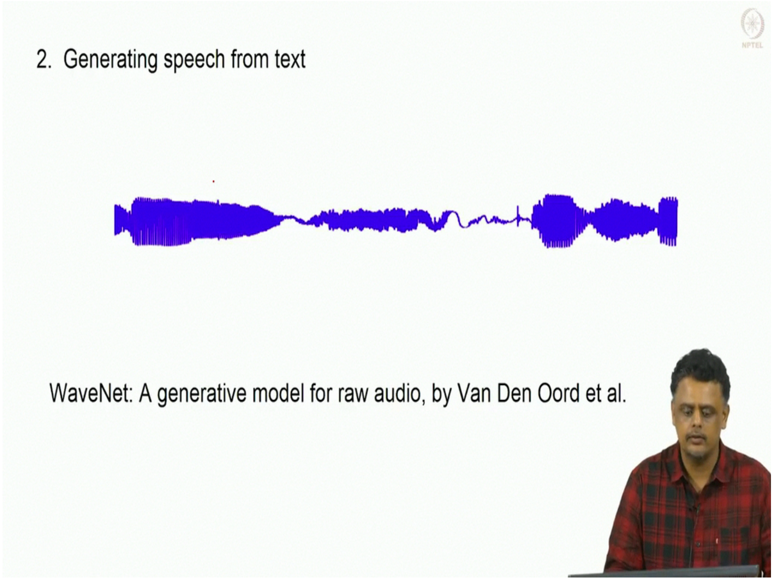

So, but if you, but the more deeper applications, for instance, generating speech from text right, so if you have say programs like an automated assistant that response to your phone call, that is the kind of program that it could require, so it has some text stored in digital format and should now convert to your speech signal in order to respond to a caller right, so for raw audio okay, so you would like to have a generative model, if you think of it.

(Refer Slide Time: 9:23)

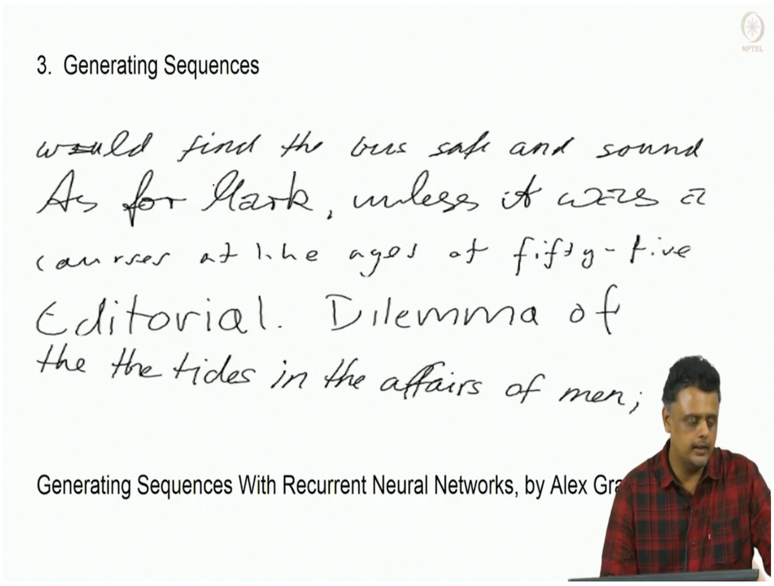

Generating sequences of text, so here is one example, but if you can think of it, a simple application, you know it is your auto complete in your cell phones or when you writing mail programs that you are using might have that, so if you type a few words, it should be able to complete the word or you may be even the sentence, that is also a generative model right, so those are some of the more obvious practical applications of the generative models, okay.

(Refer Slide Time: 9:51)

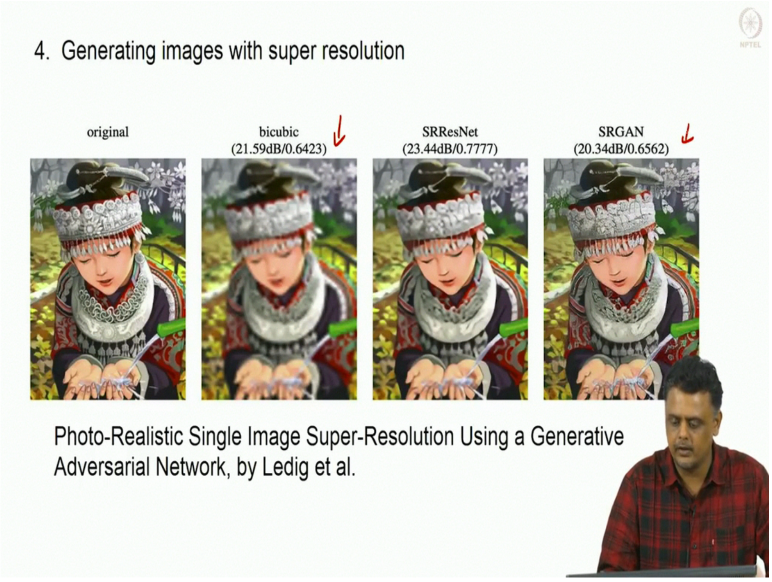

Another application would be super resolution in the sense you are given, so this can think of two different ways like, so here you are given a low resolution and ideally you want the code to generate high-resolution data of the same, you can also think of image in painting, so wherein if there are some holes in the images or some damage areas in the image, you would like to read that, so if you have a generative model, it should be able to figure out what the missing data is, so you can think of it for as a data imputation problem also right, so these are some of the, you know very high impact or also some first-order application for generative adversarial network, okay.

(Refer Slide Time: 10:41)

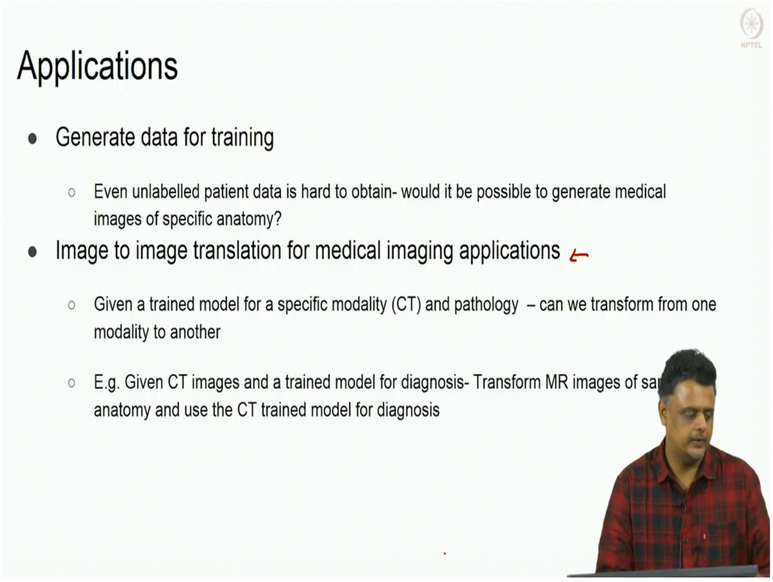

So in the context of medical imaging, this also has, it is very powerful uses, so some initiate ones that has also been reported in literature, so for medical image applications, you not only need patient data that has to be obtained with patient consent and there is a lot of, you know work that has to be done in order to make sure that you are assembling the correct patient population and all that.

So given that it is hard to obtain label patient data right, so in fact just to obtain the patient data and then of course you also have to label them. That is even more difficult problem, so that you can train a classifier but it would be nice if you can actually generate data that is generate images of anatomy, so you have, let say brain tumours right, so you want to analyse brain tumour or liver tumour, it will be nice if you can have a generative algorithm for generating images of the liver or the brain, even regular anatomy that will be nice, so that we can actually do at least anatomical level segmentation okay.

So that will be a very immediate application for generative models, there is another need application would be image to image translation, here it might, this is a little bit complicated from the, try to understand this, so in many situations, you know you will have a lot of data for a particular morality, so let say CT images of the liver. Okay, they require a few scans would be easily available, so large number of, even unlabelled data is easily available right.

So then you can train a model to whatever diagnostic tasks that you wish to complete right, let say CT images of the liver, you have used CT images of the liver to, let say segmental liver right, just to mark the anatomy of the liver, what then the new MRI scanner comes in and you are starting to take MR images of the liver, which is generally maybe, not so easily available right, so in order to train a supervisor classifier for liver segmentation from MR images of the liver you need a lot more data, which might not be available.

Then, so a generative model in this context can be used to convert or translate the MR image to CT image, use the CT image segmentation model and then transfer the segmentation mask onto the MR image, so that is a possibility, so you have spent, you have train an expensive model with a lot of data and a very similar problem comes along, it will be nice if we can use that train model where you spend a lot of time and resources doing that alright.

Of course this does not work for everything, so just to for those who are not in the medical imaging field, so you cannot train, you know the CT images of the head, let say you have network that does some diagnostic tasks on CT images of the head and then you cannot take MR images of the liver and use that network without, of course you can do, find tuning and then change, that is different transfer learning problem, we not directly use that network, there is no image to image translation here we have to use it in a correct context, so here are looking at CT images of the liver and MR images of the liver, than it is more meaningful to do that.

Of course these are slightly more research oriented applications, but very high impact they have are solved, so not only in this context, but there are various other context in which, wherein generative models are used specially in medical imaging task, where very useful, especially in scenarios where we were, you do not have too much data, this might come in handy.

(Refer Slide Time: 14:43)

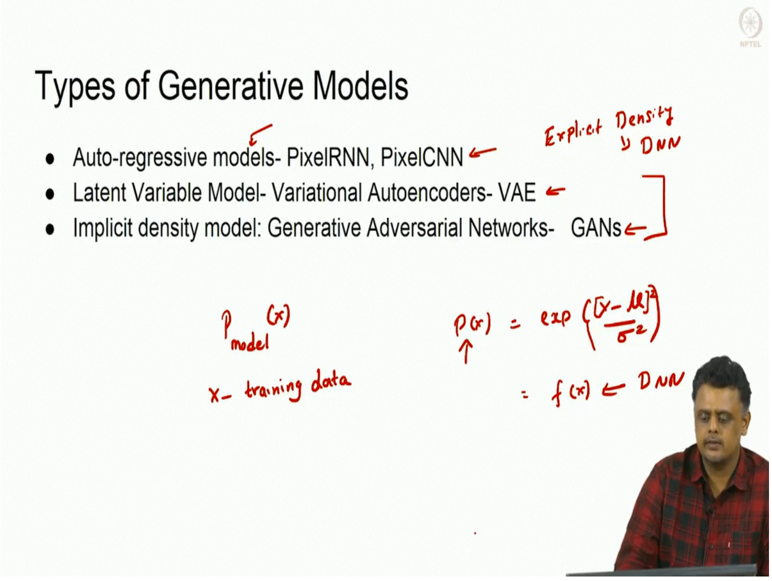

So okay, so what are we types of generative models that will cover? Right, so will go into details, couple of these techniques, not all of them and so how they can be used to generate data, so most of the applications will be looking at, would involve generating images okay, that is the most common application, so far, we will not look at text or speech, but just generating images, so that context, there are different types of generative models, so this is one of the more highly cited models PixelRNN or PixelCNN, they are called, they are based on what are called autoregressive models. Okay.

So they are called explicit, so explicitly determine the probabilistic, the probability density of the data. Okay, so there are called autoregressive models because they, if you think of a image, you try to generate a particular pixel based on, are the pixel you have looked at so far, so if you raster the image row by row or column by column, try to predict a pixel which is the second column by considering this is in the first column of the image, that is the basic principle.

So it is comes under autoregressive models and it gives you an explicit density images basically the neural network itself, the deep neural network, the densities model using the deep neural network right, so just to give you an idea, so, for instance, if you are working with some simple data, you can model that data using a let say a Gaussian right, so your density would be something of this sort okay, right something of this sort.

So instead of this, we will just have some of X which is basically nothing but deep neural network okay, so in most of the applications, most of the generative models using the neural network, the idea is to model P of X with the deep neural network itself, so the network is the model okay, so another class of generative models which also explicitly determines density is the latent variable models, one of them are more highly cited work is the variational autoencoders.

So we will cover this in the next few videos, so variational autoencoders again, they are, it is call variational autoencoders because there is a, it is an approximate way of determining the density function right, so that one and we see why it is called latent variable model when we actually look at the algorithm in detail, of course, the most studied model in the recent past is the generative adversarial networks. Okay, here we do not explicitly model the density with the probability density, it said we just sample from it direct okay, so the neural network outputs the sample right, so directly outputs the sample, so that is the idea behind GANs and again these are deep neural network and primary has been used for generating images.

So, in fact, I have not seen any application. Otherwise, but it is the more popular application or for generating images and there are being very recent successes for generating images of human faces and they look very realistic, so you can look them up if you do a search, so we will primarily focus on VAE and GANs okay, so variational autoencoders and generative adversarial networks will be the two generative models that we will focus on, so to summarise, so generative models try to model the underlying data distributions.

So data is basically your training data, so we want to figure out this, we want this, where X is a training data right, so that we can draw samples from the probability distribution which looks like you are training data, so that is how we should interpret generating models, so once you have that model and there are a lot of applications that are possible, that is one of I showed you, and they can also be used as classifiers, so it can also be used in the context of classifying, of course, it is much more straightforward to, if you have labelled data, it is much more straightforward to train the classifier, directly using a deep neural network, but it is a possibility to know that using the modern that you are figure out okay, so we will continue with this in the next few lectures, will look and various how-to encoders and generative adversarial networks. Thank you.