Lecture 22.5
Bringing it all together (the deep generative summary)

And now the entire summary of this whole deep, generative stuff that we have done. Right?

Refer slide time :(0:19)

	RBMs	VAEs	AR models	GANs
Abstraction	Yes	Yes	No	No
Generation	Yes	Yes	Yes	Yes
Compute P(X)	Intractable	Intractable	Tractable	No
Sampling	MCMC	Fast	Slow	Fast
Type of GM	Undirected	Directed	Directed	Directed
Loss	KL-divergence	KL-divergence	KL-divergence	JSD
Assumptions	X independent given z	X independent given z	None	N.A.
Samples	Bad	Ok	Good	Good (best)

Table: Comparison of Generative Models

Recent works look at combining these methods: e.g. Adversarial Autoencoders (Makhzani 2015), PixelVAE (Gulrajani 2016) and PixelGAN Autoencoders (Makhzani 2017)

So, here's the summary so, we had these four algorithms that we were interested in. So, we had RBMs variation auto-encoders, auto regressive models, in which we did need and made, will not talk about pixel RNS. Okay? So, these are the four things that we did and there were several things that we kept asking about, this model can you do abstraction can it do generation and so on. So, abstraction, how many of these models were able to do? Except for in GANs do you have any abstraction; no in fact you start from some normal, random distribution and come to the image. So, there is given an image, there is no Z. Right? You don't have Z given X in a gap. Okay? So, it cannot do abstraction, whereas RBMs and VAEs can do that, what about generation, all of them can do it. Okay ?What about P of X, if I want to compute, P of X the options that you have is, tractable interactable, not applicable, if I want to compute P of X, in the case of RBMs tractable, interactable not applicable interactable. Right? Because to compute P of X you have to sum out the edges. Right?
Remember that so, it's interactive, for variation in auto-encoders. So, remember always, you had this bad expectation and then you approximated it by something, in the case of RBMs, approximated it by, sampling in the case of variation auto-encoders, by using variational inference. Right? So, you did not optimize the actual objective, you optimized the lower bound. So, there is a difference that means it was interactive, what about auto regressive models? Tractable but slow, what about GANs? Not applicable GANs do not even compute, a P of X, what about generating samples or sampling from this distribution, in the case of RBMs, how will you do it? How will you generate samples once you have learned the distribution, you will do? Get sampling for variation auto-encoders, you can just pass a Z, I mean for variation or ten good as it was simple and you take a Z and you just generate the image from there right the decoder will do it? For AR models, you can do it, but it's very slow, because you do it one pixel at a time, from left to right and for GANs it's very fast, because you just pass it through your convolutional neural network or the transpose convolutional neural network and you'll get it. Okay? What about the type of the graphical model? The options are directed undirected, undirected directed, directed directly you take a Z and then give X. Right? It's also directed. Okay? What about the loss function? Okay? This since you have not done the math part. So, it's KL divergence for the first three and something known as Jensen's Shannon divergence for the last one, which is again a sum of two KL divergence. Okay? So, why

is it KL divergence, we were always looking at cross interpolate, then why am I writing KL divergence, because the KL divergence, is the same as cause entropy, which we are trying to remember that, you have proved this earlier, because one term depends only on the true distribution. And the other term depends on the, true and the predicted distribution. So, the first term does not matter and the second term is the same as the cross until, everyone remembers that you have derived this at some point, if not you can go back and check it what were the assumptions, made by these different models, for RBMs X's are independent given the VAEs no VAEs. We did make an assumption; it was implicit sitting somewhere in the corner.
We assume that the covariance matrix was that means the X is are independent given the Z for autoregressive models; we do not make any assumptions, for GANs not applicable because you don't really find P of X. Right? Okay? And Okay? This last part is so, the generated sample. So, what's the current state of the art? So, you have these four different ways, of generating now. Right? And doing all these applications, which I said given an image and it another image and so on. So, what's the state of the art? Right? Now so, RBMs are pretty bad at it VAEs ok the images are a blade flurry they're not very sharp, AR models are good. And GANs are considerably the best among these, four options that's what you hear a lot about gaps. Okay? So, this is one way of analyzing all these models and like finding the relative differences, from each other the other is we could actually and. Okay? So, before I move on. So, there's a recent lot of newer works, which look at a combination, of these things. So, you take the auto encoder idea, you take the adversarial idea from GANs and try to come up with a net do a serial or variation auto encoder. Right? Or you take the pixel are in an idea, you take the variation auto encoder idea and try to come up with a pixel variation auto-encoders. Right? So, these are lot of combinations of these things are now happening, why because all of them have their own relative advantages and disadvantages. So, if you combine them, hopefully the disadvantages should disappear and the advantages should multiply. Right? So, that's the hope. Okay? So, the lot of recent work happening around that.

Refer slide time :(05:45)

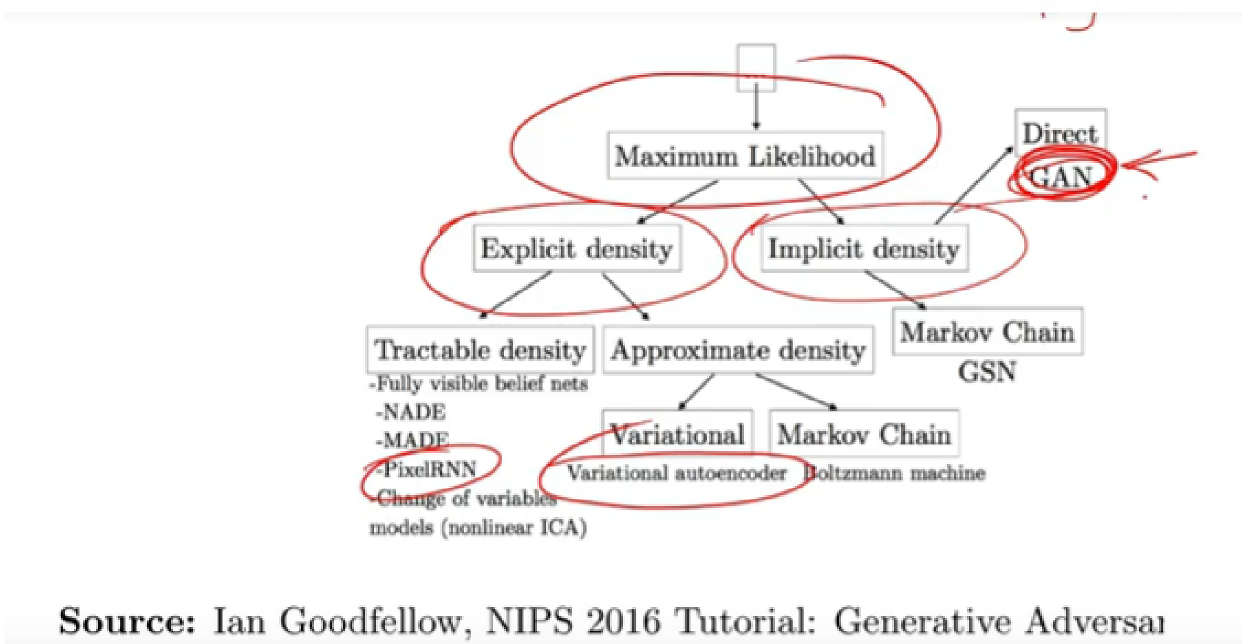

And finally if you look at the overall taxonomy of these approaches. So, remember this maximum likelihood, this nothing this is nothing but the recipe which I told you. Right? Whenever you want to learn a joint distribution, what you do you? Have a P of X and you want to maximize log of P of X. So, all the models that we do, have the same recipe, they want to maximize some log of P of X. Right? But some models have an explicit representation for P of X; some have an implicit representation for P of X. So, the explicit representation ones are GANs that's the most important one. So, the implicit representation it does not compute P of X, it just gives you samples from the distribution. So, that's guys for the explicit once, you have two parts one is tractable models. And the other is intractable models. So, the tractable models are made, made and pixel RNN, because you can actually compute P of X as the default factorization. Okay? And the interactable, ones are the ones whose uses approximate density, are variation auto-encoders and Boltzmann machines or restricted Boltzmann machines. So, this is the overall taxonomy, of all the algorithm that we have seen and these three. Right? GANs pixel RANs and variation auto-encoders are the most popular ones that you'll see nowadays. Okay? So, Okay? So, with that I officially, end the course and I hope you enjoyed it, I really enjoyed teaching it, I hope you also enjoyed learning, from it and picked up a few things along the way. Now I would like to do an important thing, I would like to call all my tears on stage and really thank them for all their support and effort throughout, the course and they've done a very, very fantastic job, it's not very easy to make these slides I and I can really be a pain in the neck when it comes to these things, replaced Nick by any other word that you want. And so, it's very, very hard because I make this storyboard and I am very particular about that, this image has to come at this pixel location and not anywhere else and so on. And the equations have to be in a certain point and so on. So, they've lived up to all that and survived through it. So, they really deserve a big hand of applause. So, please join me in doing that.