Introduction to Medical Imaging and Analysis Softwares
Professor Debdoot Sheet
Department of Electrical Engineering
Indian Institute of Technology Kharagpur
Module 4
Lecture No 16
Retinal Vessel Segmentation

So welcome and this week we are going to start with some of the application areas and challenge problems of practical applications of medical image analysis. So we have 5 different topics which range from ophthalmological image analysis to digital pathology applications and on the way we would also be going through radiological imaging and image analysis problems as well.

(Refer Slide Time: 1:14)

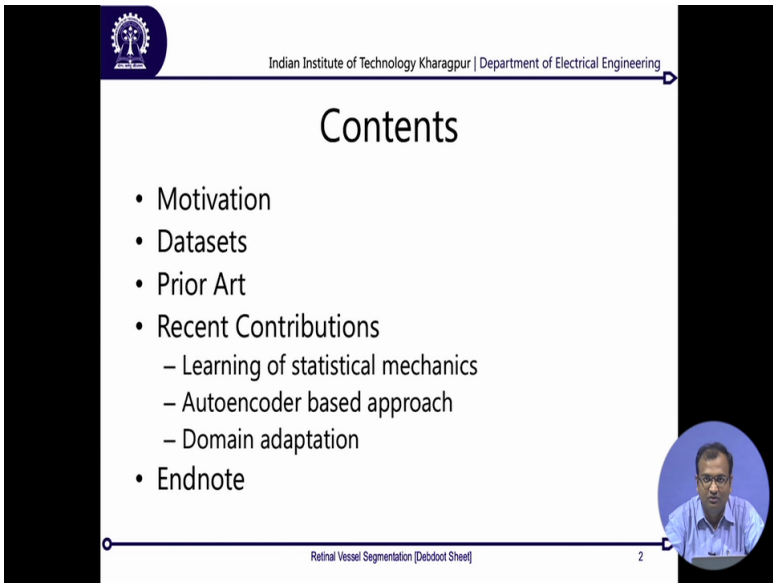

So for the start of it, today's topic is actually to do with ophthalmology and ophthalmology to make it much simpler is basically imaging of the interior surface of your eye. And we are going to look into one very specific challenging problem called as retinal vessel segmentation. So without much of a delay let us get into what this comes down. So I have this organized as the motivation, then I would be discussing about the datasets then a bit of prior art history, because this problem has been there for more than 20 years now and what has happened in the last 20 years.

And then I would specifically discussing more on the recent contributions and more of these contributions are something which comes from our group itself. And so we have three different journals of contributions coming down, one of them is by understanding the image

in physics itself and then trying to put down some sort of a machine learning algorithm on top of it in order to enhance understanding of imaging physics and taking down queues from imaging physics to do an image analysis.

The next one is a pure data driven approach, which makes use of a deep neural networks in order to do image analysis from the perspectives of doing a vessel segmentation of retinal vessels. And the third one is quite an interesting one which is where we apply this very descent concept of what is called as domain adaptation. And we show you as to what a domain is what we mean by a domain and what is about adapting a domain and a very curious case where you would see that actually trying to do it via a domain adaptation improves the accuracy with which we can actually segment out vessels.

And finally I would leave you on a end node with a point at to one of the very fundamental papers on retinal image analysis, which you are expected that you go through it so that you have a comprehensive overview of what the field of medical image analysis with respect to retinal images has been till now. So that does not restrict itself only to understanding vessel segmentation but you have everything over there, from vessel segmentation to extrusion detection to understanding of fluorescent images as well.

(Refer Slide Time: 2:58)

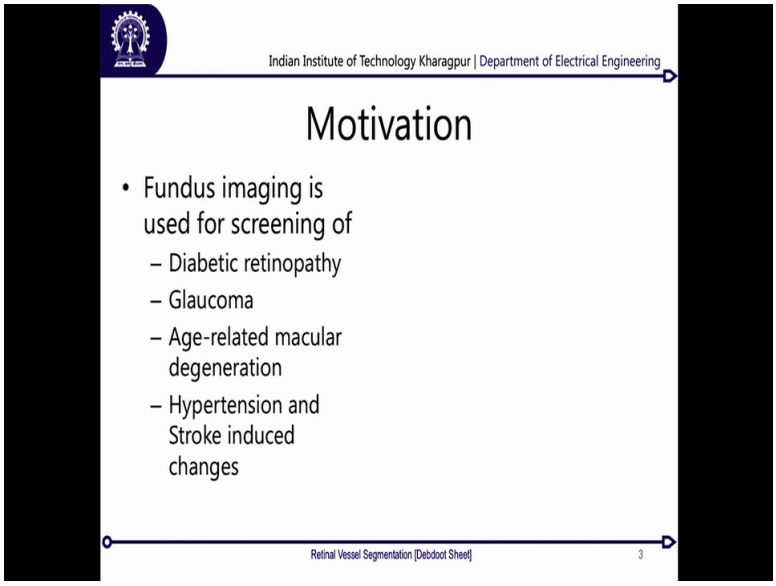

So let us start with the motivation and where it comes down. Now the motivation for this is actually to start with Fundus imaging which is primarily used for as a screening modality. Now, when you go down to say a ophthalmologist for your eye inspection, so often you would have seen that he carry something which looks like a torch light and then he would be

pointing that into your eyes and then through it looking down within the interior surface of your eye. Now from this whole thing what he gets down is image of the interior surface of your eye which is the photosensory layer which is called as the retina and the soul surface together is called as fundus and that is why you get this name which is called fundus imaging, so you are imaging fundus over there from there it comes the name.

Now this is a very good indicator for diabetic retinopathy, glaucoma related problems, then you have age-related macular degeneration, which is the macular or the photosensory surface the layer between the photosensory surface retina and then the other underline surfaces is where with ageing you would have some sort of a degeneration coming and it starts to get regally and curved and you would be having problem with your vision.

So from there to your eyes actually have a good early indicators for stroke and hypertensive changes as well, which is the reason why image analysis for the retina is now really getting a high push in the community. So these challenges over here are detecting and segmenting vessels, fovea, optic disc, so we have on the optic disc side you have both localization and segmentation as the problem. Then understanding pathology and detecting pathology from your fundus images as well. The next part is obviously on image quality assessment, so this is about how good is the quality of these retinal images which have been acquired by the ophthalmologist directly.

(Refer Slide Time: 4:54)

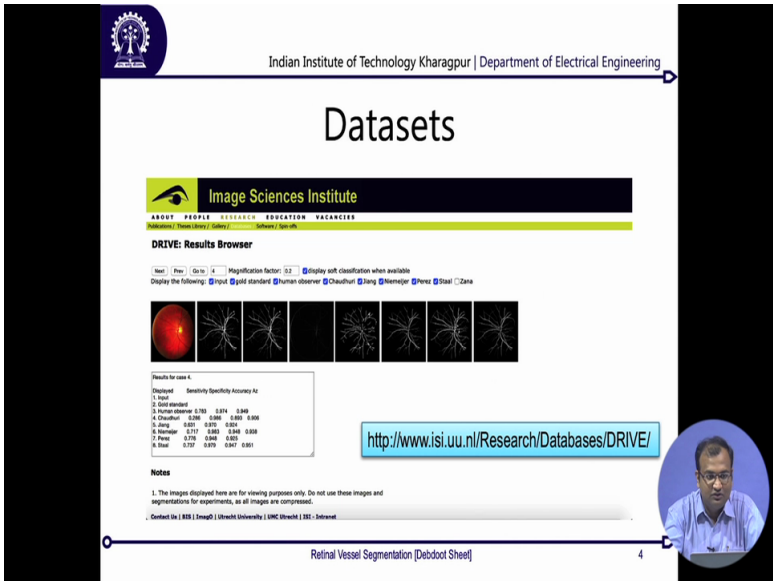

Now from there I would take you down to where is the play field of this one, so one of the most popularly used widely employed one is what is called as drive or the digital retinal

vessel extraction dataset. Now these ones are available for free use, you can just go on to this particular website which is on the pointer over here. Now what you would see typically is that I have opened up as a gallery, but you can always go to the downloads page and then download all the 40 images, which come as 20 for training and 20 for testing sets over there.

Now, a lot of people who had contributed earlier in terms of prior art, they have put back the results over here and then you have the results browser, where you can click on one of these and then it would open up and show you each of these images over there and you can see the accuracy, sensitivity, specificity and area under ROC curve these values coming down over there. So this is one of the main playgrounds which we used for vessel segmentation.

(Refer Slide Time: 5:52)

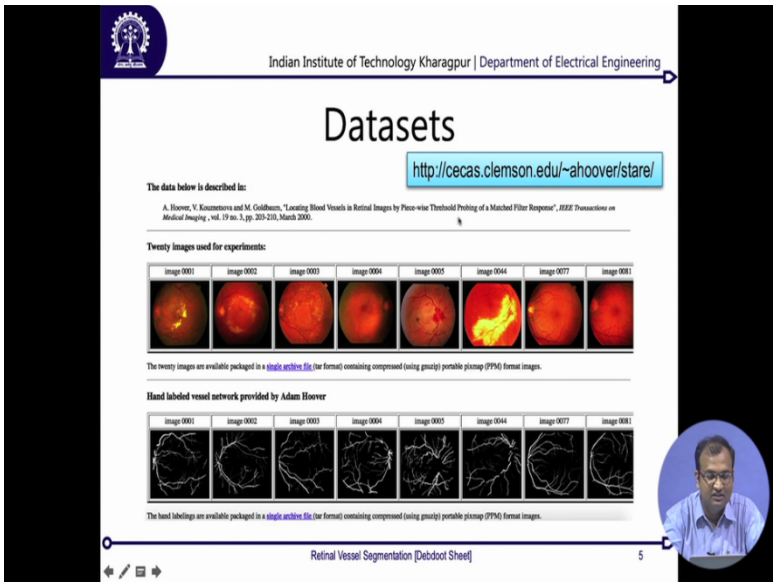

The next playground which we do is what has come up in the recent years is called as stare and the beauty is that while drive is primarily or this previous dataset drive, this is primarily from healthy subjects who are not symptomatic of any particular kind of a retinal pathology. The stare dataset is from people who have some sort of a pathology over there and for most of them they do have some sort of a hypersensitive disorder or they have diabetic retinopathy problems over there.

So these have both healthy subject data as well as disease subject data and since we had already discussed about systematic evaluation or validation and I had explained you about the ratio of subjects from different categories and why you should be preserving the priory probability of for each of them. So now you have a very clear understanding as to why we need to have both the symptomatic cases and asymptomatic cases coming down or disease

cases and the healthy cases their within our dataset in order to evaluate the performance over all algorithm, so this is one I would definitely put down a pointer, there are a few of the other ones

(Refer Slide Time: 6:56)

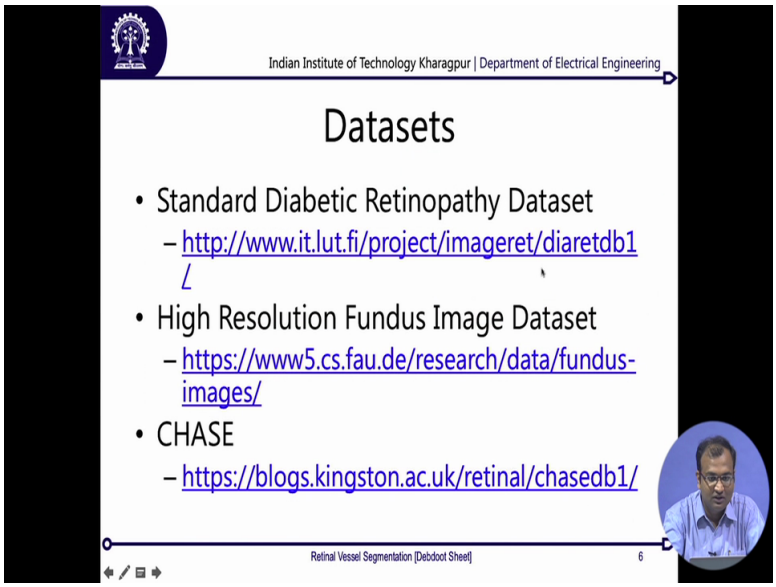

So one of them is direct which is a standard diabetic retinopathy dataset and there is a db0 and db1 both of them you can just make use of them, they come down with ground truth annotations as well so you do not need to worry a lot about it. Then there are high resolution fundus images, which are from the University of Aalen at Germany. And these are good in a way because they are full HD scale images, which provide you a much better resolution and much higher granularity then you have for other images which were sort of in a 600x800 or a 640x480 resolutions image pixel sizes. And then there is another one from the Kingston University in UK, which is called as the chased, so this is come up in the recent years only and you can make use of any of them, they are all free for download.

(Refer Slide Time: 7:52)

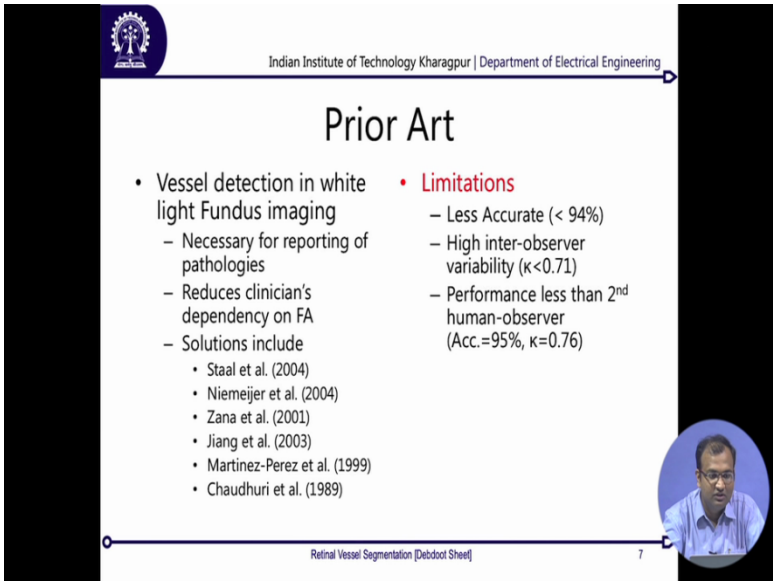

Now that you know about all your datasets, the next thing comes down is knowing about what has happened in the past and that brings us to the prior art problems over here. Now vessel detection in this white light fundus imaging and the major necessity is that you have these for the referencing of different pathologies. So whenever, a ophthalmologist reports a particular pathological manifestation on the fundus image itself, he or she is going to report it with respect to their anatomical locations or landmarks which are often actually denoted with respect to vascular branching or where the vessels are, so that is what access a ground referencing.

So say as if you want to tell our address in say some sort of a remote village where there are no street names and street numberings or house numberings possible, then generally you have a referencing over there. So it is at the cross section of so and so road where you have this big banyan tree and then you need to look at the second house on the right. So, these ophthalmologist would also be using a similar kind of thing because obviously vessels in your eyes are not numbered and all pixel locations over there also something which is not associated anatomically by a number.

So you need to put down this kind of relative referencing in order to come down to reporting out these pathologies. Next is that once we have this detection of vessels done quite affectively, it would reduce significantly clinicians dependency on the fluorescent angiography in order to find out how the vessels go. Otherwise, the procedure would be that

in order to find out the vessel map I will have to do some sort of a procedure which is called as angiography.

That would mean for this particular purpose what clinicians do is they inject a dye which is fluorescent dye. Now as this dye passes on through your blood vessels they keep on imaging your blood vessels and then they integrate over time. So the dye is going to pass through your blood vessels directly and if I am integrating it over time, then what I would get down is basically a map over the regions from where this dye has passed down.

Since it is passing through blood vessels, so you get a map of all the blood vessels coming down, the problem is that this dye is generally toxic to our body and you would need to flush it out of your body. And so people would need to pass down more of urine, drink lot of water so that it gets filtered out of your kidney and then flushed out of the blood and generally these kind of imaging is done for people who are diabetic and who have a diabetic retinopathy problems, so that you have a vessel map available. Now people who already have diabetes they have a kidney overload already there and you are going to overload them. So this is sort of a quite risky situation on the clinical side as well and having a automated system for detecting vessels on bright white light images would make it much more simpler.

Now you till now since we have not seen that many images you would often lead down to this confabulation may be it is not so complicated that you would have to put down a computation technique over here, but I would come down to the exact problems as to why we have a computational technique because by bare eyes we would often failed to find out those vessels. So once we go down to those images I would be showing you pin pointed spots where that particular ambiguity comes. Now since it has been active for the last 20 years, so you can look down that these were the initial contributions which were there and a lot of them basically made use of derivate based operators or filtered out derivate. So there were laplacian of gaussian hessian kind of operators which were used over there.

And on top of that people did relay on texture measures as well in order to find out, but most of it was related around one basic concept that these vessels they are basically what appear as lines over there, so if I have some sort of filter kernels which can emphasize on these line kind of behaviors, not necessarily straight lines but some sort of a contour which is going down through the image, then I would be able to get down my vessel very easily. Now, based on this all of these methods which have been done, now we are specifically not going to discuss about each of these methods in detail over here but rather look at what are the recent

contributions which have been able to overcome the challenges which were still posed by these particular methods.

Now, what I do over here is I have a basic listing of what that means. Now, generally all of these methods they are less accurate, so it is below 94 percent, the reason from where this comes down is that what these dataset developers had done was, they had used two (()) (12:17) observers to manually annotate it. Now it takes about one hour of time to manually annotate all the vessels present in one fundus image over there.

Now, when you take two different observers in order to do that what you would do is in essentiality you can take one of them as a ground truth and the other one is who is getting evaluated and then you can do a inter-observer variability study. Now, if we look at the accuracy between two observers when being compared, so that is what comes down to 94 percent and all of these methods which you see over there none of them had an accuracy which was rivaling at least 94 percent, so their methods were lower than that.

Now, this obviously brings to a point that my computer is not as intelligent as humans so that I can actually bring a computer in order to do this kind of a repetitive task over there. So for a repetitive task also it is not efficient as of now, so that is the reason why these methods cannot be employed in that good way. The other point is that there is a high amount of inter-observer variability when compared to these methods.

So and that is about like if there is a very low amount of inter-observer variability, you will have this kappa score basically go down to 1. So that would mean that there is a high amount of inter-observer consistency between them. Whereas, whenever this kappa value over here is much lower than 1, and then you have a low. So this 0.71 is what is the consistency between two human observers, so that is the bench mark which we need to always break, till we are below that methods are not sufficient to go down as automated processes.

(Refer Slide Time: 14:18)

And the other point is that none of these methods had a performance which was rivaling of the second human observer and that is why these were not taken into consideration. And from this prospective is where I am going to bring to you the recent contributions which have, so I have three basic contributions. So the first one is from the learning of statistical mechanics and this is one which we had from ISBI 2013 as one of our contributions over there.

(Refer Slide Time: 14:20)

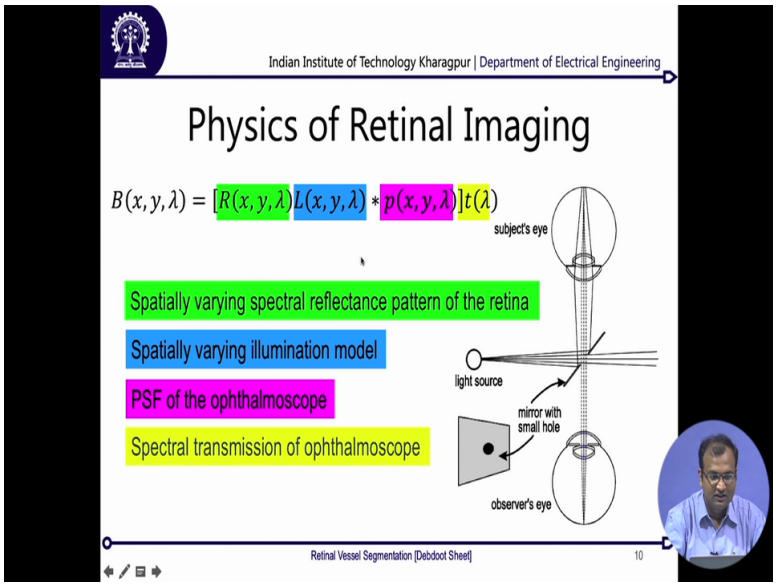

Now where this goes down is we start with the physics of retinal imaging and how it happens is say typically this is the eye of the person whose eyes are being investigated or what is called as a subjects eye, ok. Now typically for a ophthalmoscope what would happen is there

is a light source and the light source goes through this sort of a mirror and irradiates the eye and now you would have some reflected light coming out back from the eye.

Now this light has to come back from here and instead of getting reflected over here because then I cannot see anything it has to go and meet down this observer's eye over here. So I have the observer looking over here through this small pin whole. Now the concept over here is that you have this light source which is irradiating the subject's eye, but does not interfere with the observers view in any way and you just see the observer can just look into what is there within the eye.

So this is the basic fundamental operating principle for any kind of an ophthalmoscope which is being used, so including that small handle one which looks like a torch light to even surgical ophthalmoscopes as well. Now, if we take all of this into consideration so there are some sort of optics involved over here, so there may be some optical transfer functions associated with this mirror then with the lenses of your own eyes and then aqueous humor and vitreous humor which is present inside over there and all of these have to create some sort of a mathematical model coming down.

Now, instead of this observers eye we can put down one single electronic sensor over here CCT sensor or CMOS sensor and then I will have to look into the response of that CMOS sensor. Now, looking into that prospective what would come down is this sort of an equation. And let us look into what the different aspects of this one would be. So this R is a contribution which is called the spatially varying spectral reflectance pattern of the retina or what is created over here.

So this is about like given I have a light of wavelength lambda and it spread in a location of x, y, so there can be in-homogeneity in the incident light or the reflection pattern, then it is going to have a response created by this. Now the light which is incident over there that is also a spatially varying commodity and that also has a contribution of the wavelength lambda and it varies a long x and y, so that is called as x that is call by L over here.

Now this whole thing is what will have just a multiplicative nature over there because 1 pixel to 1 pixel is the nature. Now along with that there would be a point spread function of this ophthalmoscope optics over here and that would in a sense mean the convolution of this PSF function with whatever is coming back from the subjects eye. And so we have this convolving operate which does a convolution and this is what comes out.

The next part is that there this ophthalmoscope over here will also have some sort of a spectral transmission pattern or what is called as a spectral filtering approaches over there, so it will allow certain wavelengths to go and it stop certain wavelengths over there or it can have a different kind of amplification factors for different wavelengths. So for some wavelengths it will do a all pass, for certain wavelengths it will restrict it to say 50 percents for some to 70 percent and it will allow it to go.

So together this is what will give me the model called as B(x,y,lambda) which is the total response being received at the sensor electronic sensor placed over here. Now, from there let us go it down into the understanding about what happens at the digital read out when we have on a electronic sensor. So say we have sensor which has a bare pattern and let us consider a standard bare pattern over here on as in this check matrix.

(Refer Slide Time: 18:00)

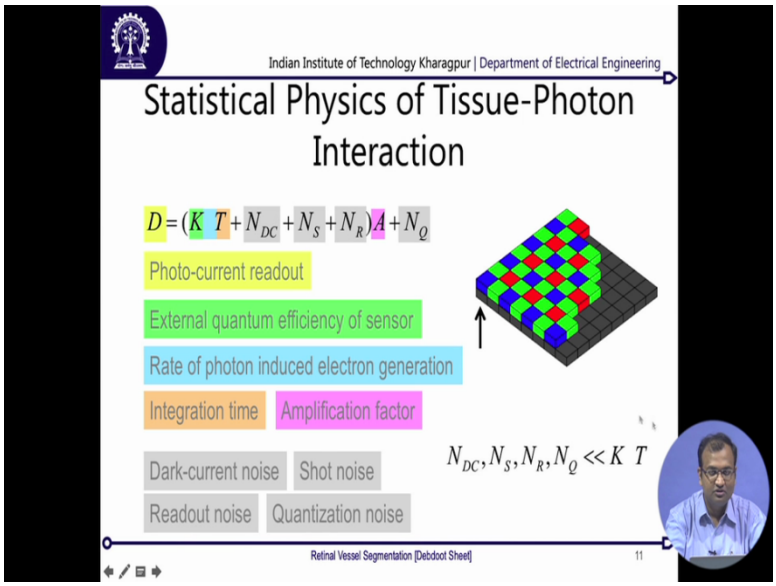

Now you would have some sort of a current which is being read out at this particular at any pixel location and that is called as a D. Now that will have some constituents of noises as well and these noises will be the dark current noise, there will be short noise, there will be read out noise and for much more details you can actually refer to some electronic sensing books, which will have much clarity on this one. So this would what will come down as a very basic model of photon to electron paired generation and then digital read out sensing from a electronic sensor. Now, on this kind of a system apart from the noises you would have a external quantum efficiency of the sensor. Then there would be a term which would be the rate of photons induced electron generation row, then an integration time T and you would have an amplification factor A.

And together in any kind of a general operating condition what would happen is that this noises would be much lesser than K Rho T over here, which is the rate at which photons are being generated and then your distribution of this current this read out voltage D over here will be what will be guided by the distribution of this K Rho T and we need to understand that part.

(Refer Slide Time: 19:16)

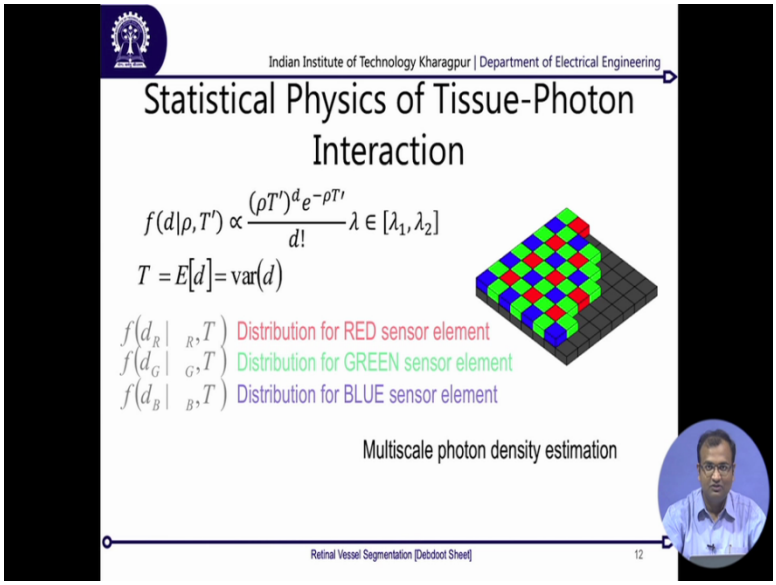

Now in order to understand that part what will happen is that generally if row is the rate at which these photons are incident on a sensor and T dash is say the total integration time of the sensor, then this distribution model say f of the digital read out data d is what is proportional to what we have over here and if you inspect this one, so this is what looks very similar to a Poisson distribution and additional factor which we have is that this is guided down for when lambda is in the range of lambda 1 to lambda 2 and that comes down from the factor that we have these sensor over here, so each is fit down with one specific kind of an optical filter either a green filter, or a red filter, or a blue filter.

Now they will have some minimum wavelength to maximum wavelength which is the range in which they will pass down all photons over there and that is what comes down over here as lambda 1 and lambda 2 specifically. So with this we know that the distribution of all the photons over here is going to be Poisson. Now given that it is a Poisson distribution process, the main objective over here is that we define a meta variable which is called as T, then the part which we take over here is that if this is a Poisson distributed process then the expectation or the mean of the variables over there is equal to the variance of the variable and that is what we take as one of the queues for solving our problem.

Now from there we can find out a distribution for the red sensor, green sensor and blue sensor independently, if you are looking at every channel of the image, Now that would help us in creating what is called as what is called as a multi scale photon density estimation model or what typically happens is since you are solving some sort of an estimator over here in this particular line, which is estimating or finding out what is the mean value over there.

So you need to understand whether you are computing mean over three samples or you are computing mean over five samples or seven samples. So if you are looking into a 2D space it can be a 3 cross 3, which on which we can compute out what is the mean, it can be 5 cross 5, it can be 7 cross 7. Now instead of trying to optimize is to which is my best scale at which I can compute, I can actually take a whole stack of multiple of these scales and over this pyramid I can compute out what will be my expectation value coming down and that is what we call as the multi-scale photon density estimation model for each specific wavelength.

(Refer Slide Time: 21:45)

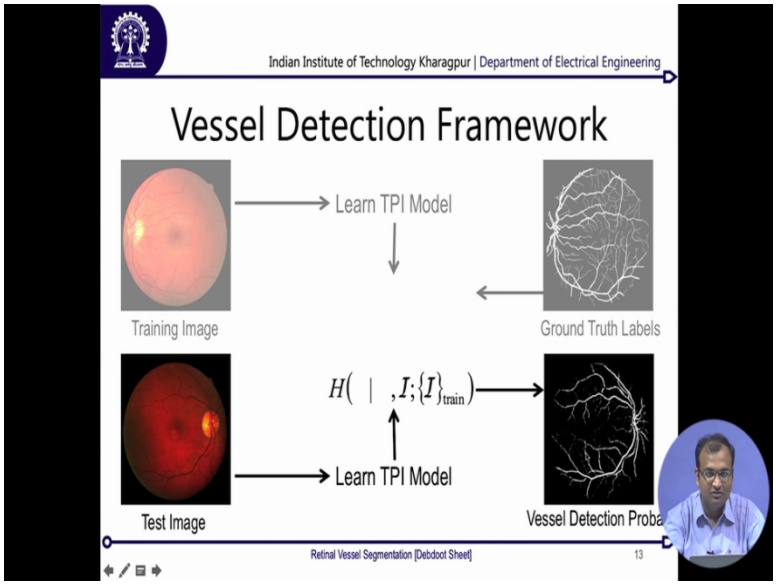

Now from there we enter into what we do for the vessel detection framework. So here given that you have a training image, what you would initially do is you would create this multi-scale estimation framework which is as such an unsupervised learning problem. So just have an estimator running and there is nothing specific which requires a supervise learning over there. Now from there what we do is, on the other side we have ground truth labels available to us, which are all of this white pixels are what denotes the vessels coming down. Now, we can use these parameters, so say a vector of all those estimations of means and then use some sort of a learning engine. So this can be a random forest, this can be neural network; this can

be a say any other thing support vector machines or linear discriminate analysis or a base learner.

And then this can actually learn to discriminate based on these tissue photon interaction models as to what and where all these pixels which denote a vessel are located. Now during test what we do is we take a test image and then we again estimate this test photon interaction model over there. Now we can use this train model coming down over here in order to predict out what will be the kind of vessel appearance models coming down. So this is what typically would be a vessel appearance model which is predicted. Now from there let us look into what this performance assessment looks like.

(Refer Slide Time: 23:04)

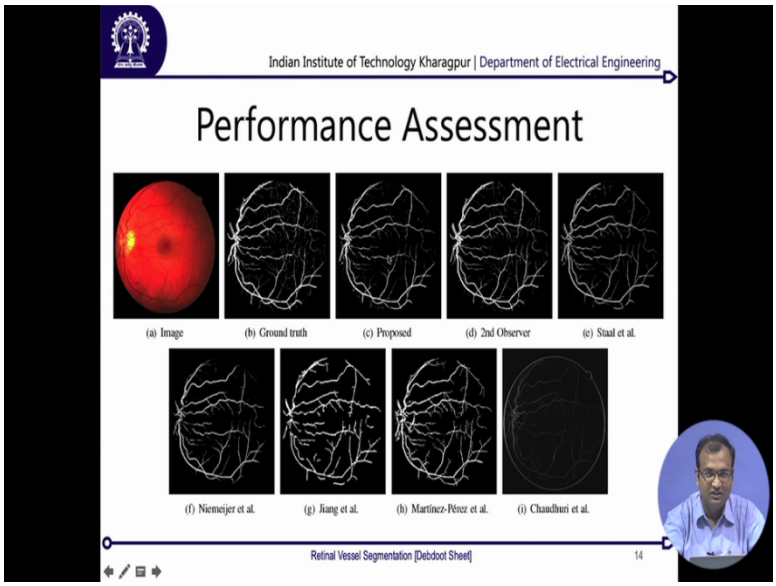

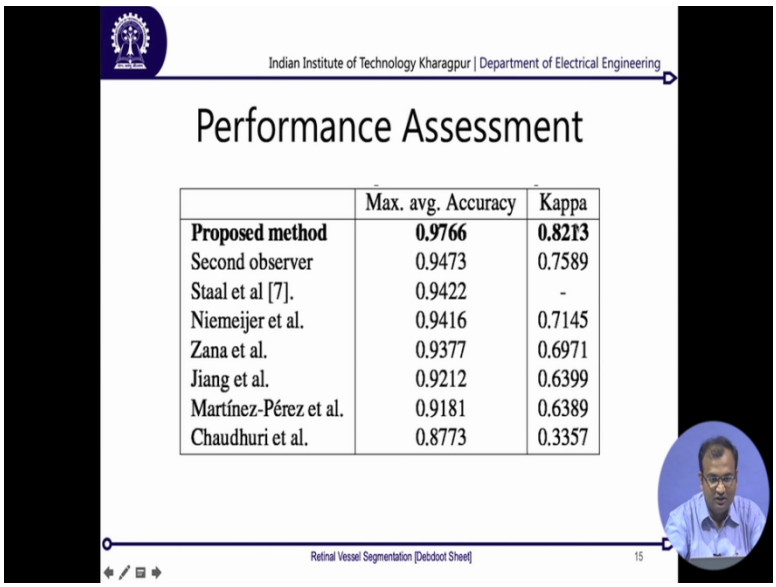

	Max. avg. Accuracy	Kappa
Proposed method	**0.9766**	**0.8213**
Second observer	0.9473	0.7589
Staal et al [7].	0.9422	-
Niemeijer et al.	0.9416	0.7145
Zana et al.	0.9377	0.6971
Jiang et al.	0.9212	0.6399
Martínez-Pérez et al.	0.9181	0.6389
Chaudhuri et al.	0.8773	0.3357

So this is what visually the method which I was describing over here looks like and they are these are all the other comparing methods which we have in place. Now in order to look into a much more intuitive numerical aspect over here, so I have so this particular method has an average accuracy which is about 0.97. Now consider this that we are looking at beating down the human observer, so the observer's accuracy, so these are methods which are now able to overcome the observers. So two humans the amount of dis-ambiguity between them we are above that. So it means that we are much more consistent than two humans are between them and this is where automation is going to actually help in the whole diagnostic process. So if you look at the kappa score as well that also indicates the same thing. Now coming down to a much critical evaluation of the different stages, so this is what I have over here.

(Refer Slide Time: 23:59)

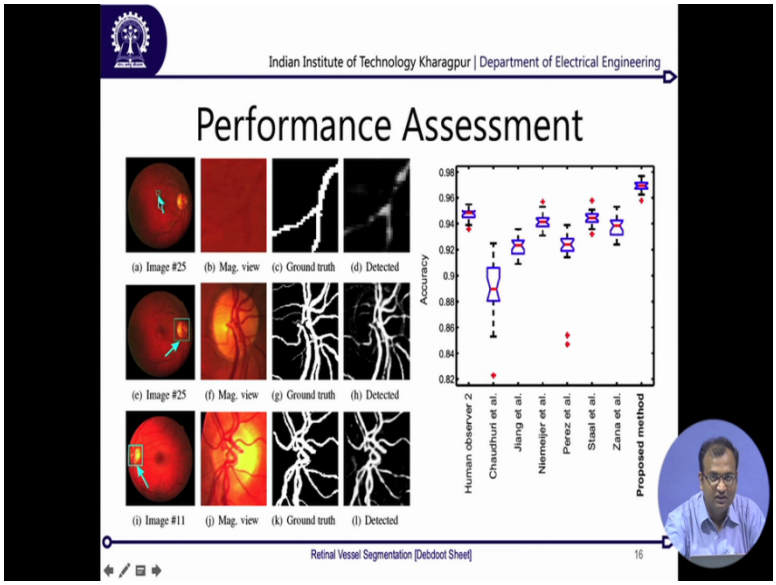

Now if you look over here, this is the main image on which this small signed colored square is where we point down the zone which is magnified and shown over here. Now if you see on this it is really hard to find out where the vessel is and then since I know where it is, so it is somewhere over here. The ground truth also marks it somewhere over here, but just looking with a (())(24:15) eye and then there is a high chance that you would miss down where the vessel is located in this particular small region. Whereas, using this kind of methods you are able to predict it out quite affectively.

The other problem is around this optic disk region, where you have different layers and overlapping vessels coming down and you can see that it gets pretty discretely predicted. The other challenge which comes down is the ambient illumination intensity around this optic

disk is pretty different for different images and that is this comparison over here; you still have the method doing it down pretty fine.

Now if you look at this notch box plot of accuracy is shown over here, one major observation is that while the average accuracy is obviously higher than all the other methods, the consistency is also higher because the spread of the notch box plot is much lower than we have for other methods over there. So this is what makes this method much more intuitive and easy to use, then where we have a major impact coming down.

(Refer Slide Time: 25:36)

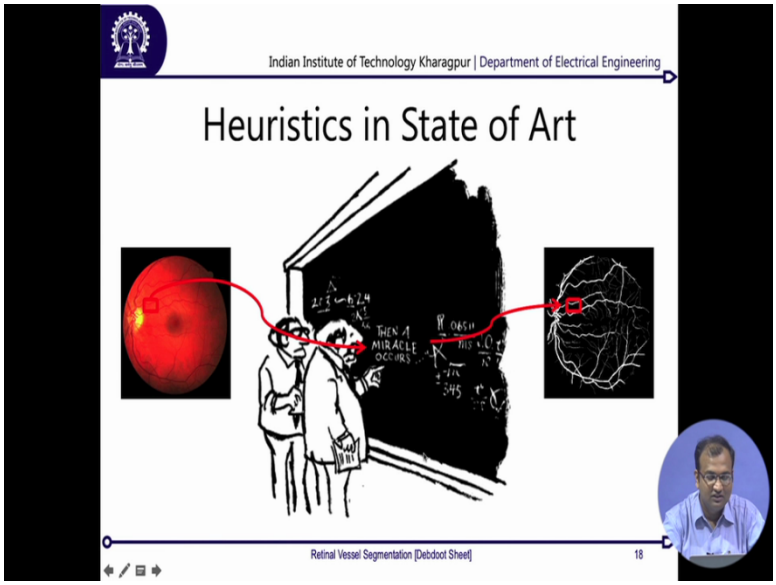

The next one is that once we have already studied about that the next is about using auto encoders for retinal vessel segmentation and this is one of a paper from EMBC 2015, where what we try to do was that if you look at the earlier method then it somehow appears like there is some sort of heuristic (()) (25:36) be going on and out of all of this mathematics you are able to prove it out.

Refer Slide Time: 25:54)

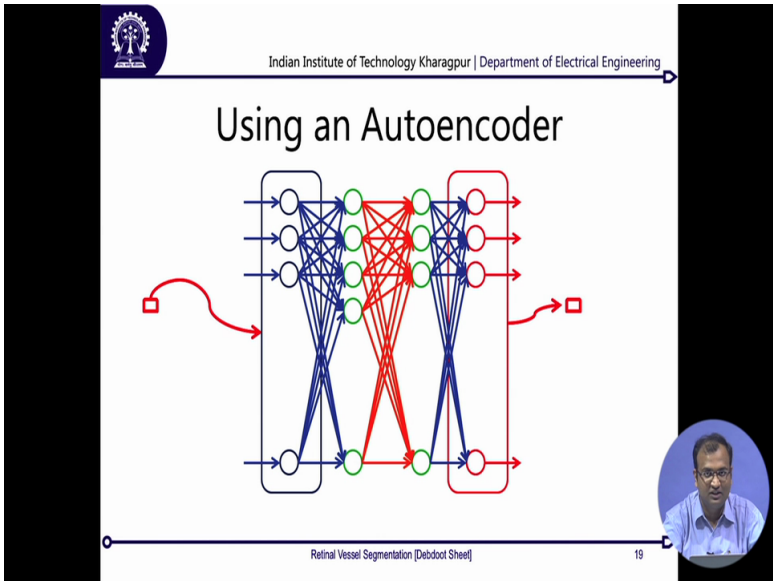

Now if you want to do it say a data driven approach, then can we do it and that is the question which we were asking over here. So the idea was pretty simple, we thought of creating just a two layer stack denoising auto encoder, so the techniques you have already studied in the earlier lectures over there and the idea was that you have similar kind of patches created over here by random sampling and then you train it with the ground truth such that this neural network can actually predict it down.

(Refer Slide Time: 26:14)

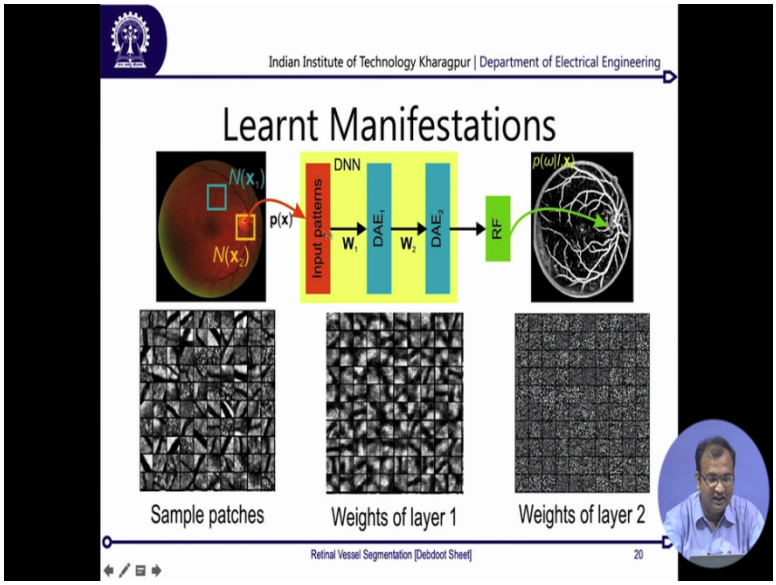

Now from there what we learnt was something like this, so what we did typically is this is the architecture which we are looking down so there were two layers of denoising auto encoders

and then there was a finally a random forest which was learning all of these layer specific outputs in order to predict out vessels.

Now here the auto encoder was used purely from a perspective of discovering features which you can learn down in order to isolate blood vessels. Now what we did is we took some random samples of patches, so these are the kind of patches which you see over here and then these were the kind of weights it had learn on the first layer of the auto encoder and these are the kind of weights it was learning on the second layer of the auto encoder.

(Refer Slide Time: 26:52)

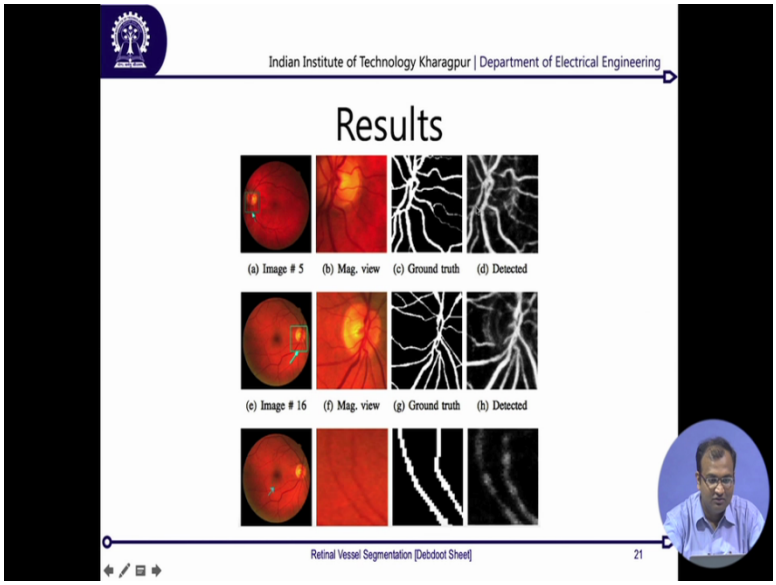

So using all of these what we get down is something a result over here, so not exactly on the same scale of comparison we did not use the same images, it was a different group of images which I show you over here. But you can carefully see that even these thin vessels one single pixel thin vessels they also get pretty much easily detected over here. So you can go through the rest of the paper for much more detailed about comparative studies with respect to the accuracy and kappa score and sensitivity, specificity.

So now this is where I was still speaking about only using one single dataset coming down, so that would mean that I have one single hospital where people almost have the same cohort group come down and I have the same ophthalmoscope being used for imaging every single person, ok. What happens if say suddenly I change my ophthalmoscope, the whole nature of imaging is going to change over there. The resolution of imaging is again going to change over there, and then can I use the same method for doing it without having it to retrain it with a new dataset, that is where the challenge comes down today. So that is what is addressed in

this particular paper which is on domain adaption for vessel segmentation and curiously at the end of it I will show that why there is a much more powerful implication of using domain adaption for this kinds of problems as well.

(Refer Slide Time: 28:13)

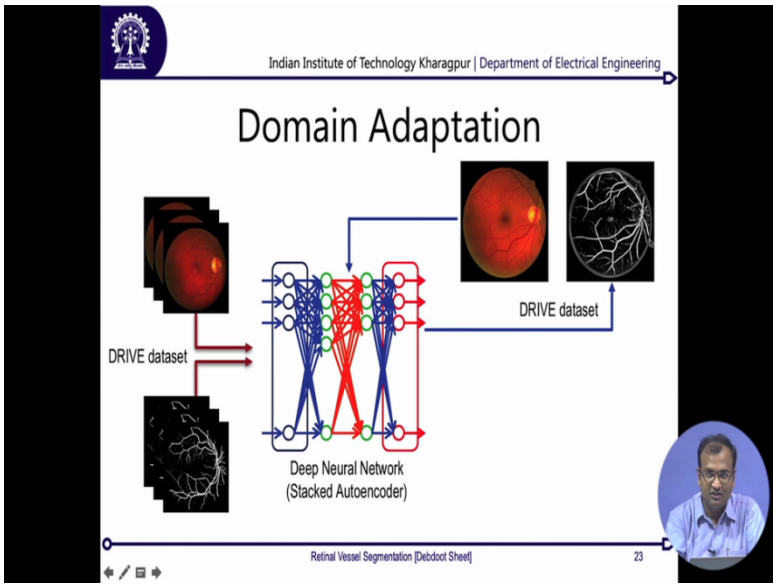

So what we start with is say I have this commonly available dataset which is my drive dataset coming down and then I train down a deep neural network, ok. So we trained down a simple stacked denoising auto encoder over here without that random forest thing over there, so this is a fully connected neural network for segmentation. Now from that what I do is, say I put down an image from the drive dataset I get a very good segmentation coming down. The moment I take a different dataset a stare dataset and do it, I have a complete change of my word. The problem over here is that all of this background which is supposed to be black now starts showing as white as the vessel along with these vessels as well.

So this is nowhere close to what I was expecting to get down in anyway and this is the major challenge which we face over here. So what we can do is something of this sort that say we have these images coming down, this is what is called as the source domain and then we train some sort of a classifier which is a classifier which is trained on the source domain, ok.

(Refer Slide Time: 29:17)

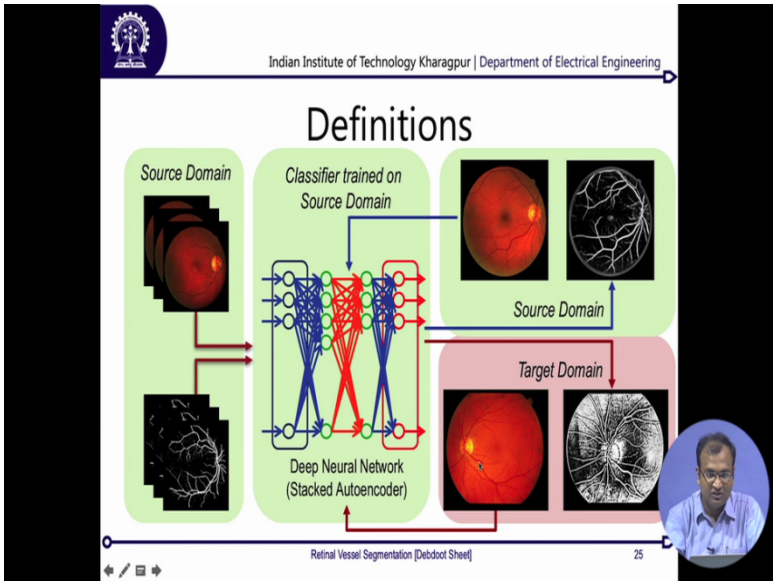

Now, we take down if we have images on testing from the same domain then we get down a good result. But the problem is that the moment we take down a different domain and images on a different domain which is called as a target domain then we do not get any of this. So this is what demystifies what is defined as a source domain and a target domain. So from there the idea is the moment you would like to do an adaptation what we have to do is we have to take image and the ground truth of the target domain and then feed all of this information via some sort of a adaptation algorithm, which is again another one level up of a learning algorithm itself into this network such that it gets adapted. And once it is adapted then you would see that you would get perfect results coming down on your segmentations in the target domain as well.

(Refer Slide Time: 30:14)

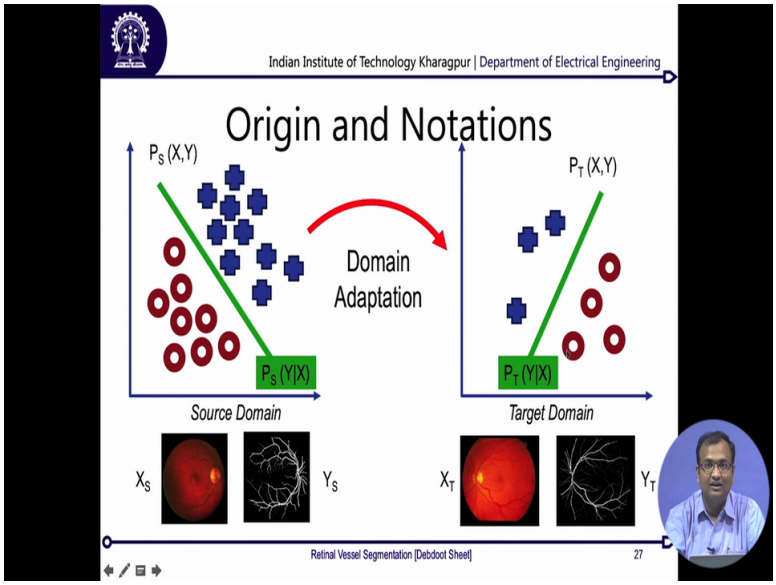

Now on the notation side of it is what we call as say these are the two classes, so may be these circles are what represents the background and pluses are what represent the vessels over there. And any kind of a classification margin is basically this line which you can draw in order to segregate them, so this is what a typical classification problem would be defined as. Now in the target domain you would have a similar thing, domain adaptations means basic definition is that can you invert this separation margin or the classification margin in a way such that it can actually adapt to classify in the target domain as well.

(Refer Slide Time: 30:55)

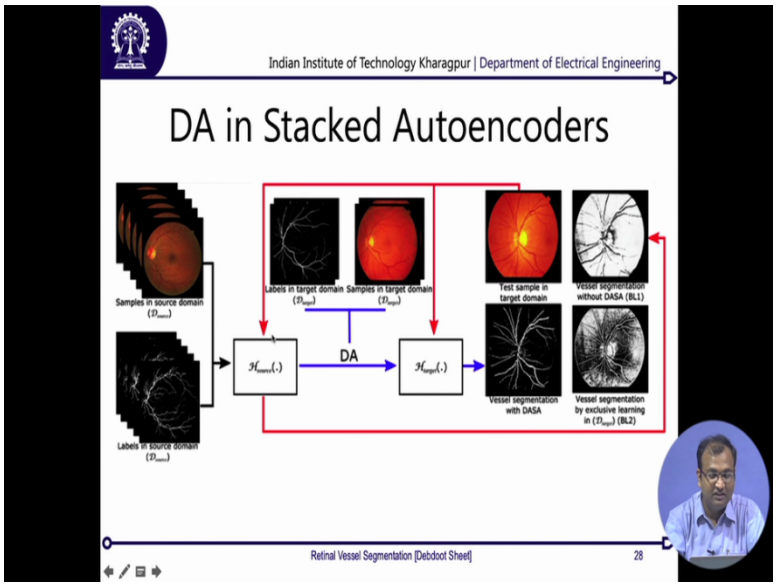

Under a constraint, that domain adaptation is where it is used likely in scenarios where you have less number of samples annotated and available in the target domain. But you have a lot

of samples available in the source domain, so what you do over there is what comes down through your domain adaptation. Now, within a stack auto encoder the basic idea what we do is by incorporating something which is called as a dropout function or what we call as a systematic dropout.

(Refer Slide Time: 31:08)

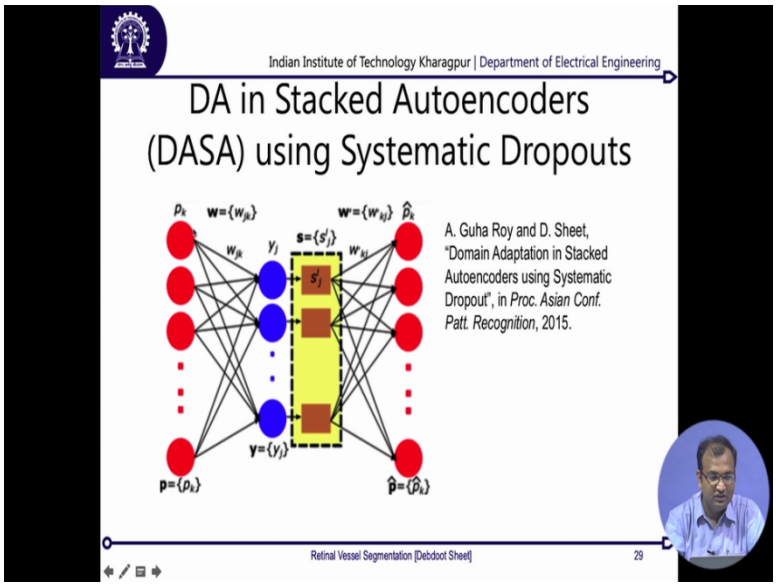

So the modification which comes down over here is something like this. And a typical auto encoder what you will have is, you have your input layers, you have your hidden layers and then you have your output layers. And then you are going to train it such that the input can be replicated as a output over there and that is the basic auto encoding principle. The problem is that if you want to adapt it, then what we do is we put down something called as switch layers in between.

So each switch is a binary switch which is connected down to these hidden layers. So either this switch will allow the response to go down or if it does not allow the response then what comes down over here is 0. So imagine it as a circuit and whether I am going allow certain current passes through that particular circuit element over there. Now this is what we have in this paper brought down over here. Now, the way in which it works out is that initially when you train on the source domain you do not have this switches available, when you want to adapt it you will place this extra layer of switches over here and then with that training data on the target domain whichever is available to you, you would be doing a feed forward and looking through what is the intensity of these responses impacting this output over there.

So those particular responses which do not impact the output over here, they are the ones which are not at all related and you can actually switch them off, so that you have a better reconstruction coming down over here or the ones which are negatively relating you can switch them off, so that is the basic concept which we have put on this particular paper. Now as a result what you get is a beautiful performance which we can see over here.

(Refer Slide Time: 32:38)

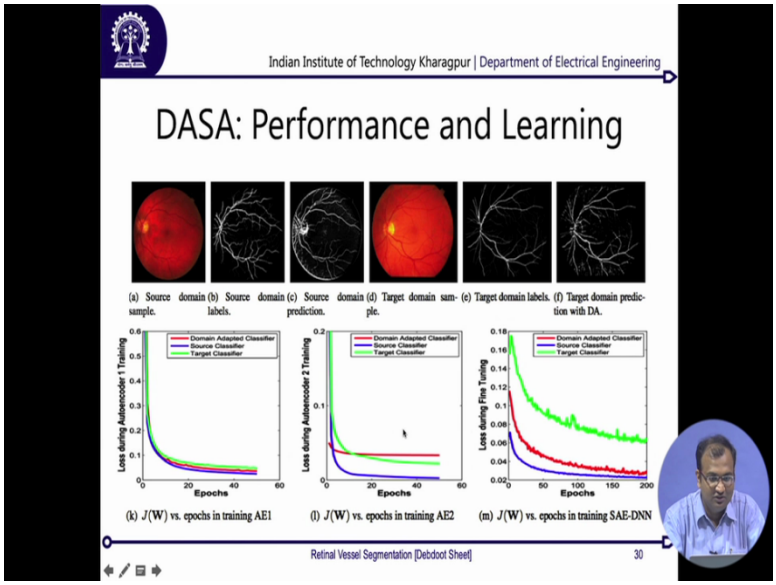

Now let us look into this final one which is the final graph of training the whole algorithm in order to get down this vessel segmentation. Now, this green curve over here is what is the result of the target classifier or basically if I use the limited amount of data, which is available on my target domain in order to discretely, train one single auto encoder base classifier. So over the epochs this is where it will saturate on the error curve, so this is a loss axis curve.

The red one is where I am going to domain adapt my classifier. So this is where, this is trained on a source domain on the drive dataset and then it is tested on the stared dataset. The blue one which you have over here is the classifier which is trained only on the source domain, so it is trained only on the drive dataset, so if I am using a classifier which is trained on the drive dataset and using it for testing on the drive dataset, this is a kind of losses I am getting down, which is one of the lowest over there because the dataset is obviously from a much less heterogeneous population and it is from all healthy people over there.

But the moment I try to train a discrete classifier on disease dataset, I see that the losses are much higher, in fact the least loss is much above the starting loss over here for healthy people. Now, in between approaches what is taken down by a domain adaptor classifier and

what you do is basically you have the classifier initially trained on your source domain which is with the drive dataset and then just use those binary switching layers of systematic dropout in between and then retrain it with limited number of samples in stare, so may be 2 samples, 3 samples just this number of images.

And then you would see that its performance is much better, and then you would do by trying to train a discrete network by using just 2 or 3 samples from stare dataset. So this is where the main beauty of domain adaptation lies, so you do not need much number of samples on a different domain to do it and now your classifier which you have already trained on a large dataset it can be used for different machines.

(Refer Slide Time: 35:06)

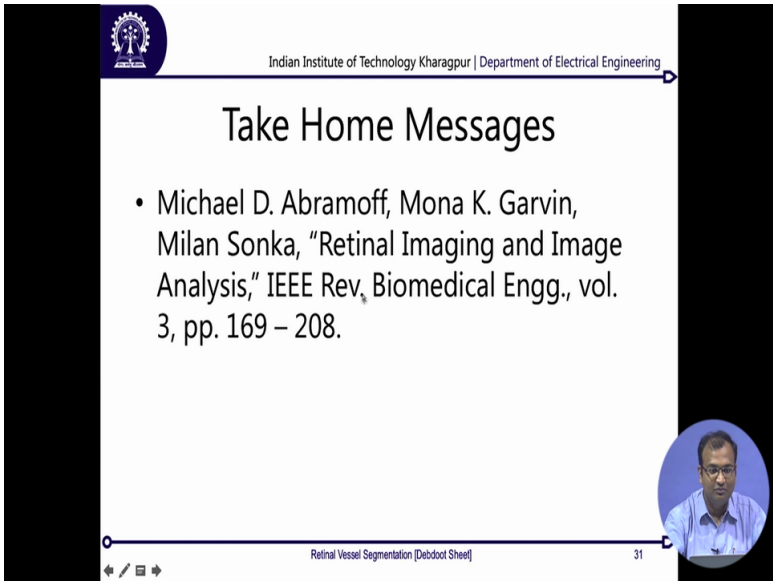

So from a software development perspective, if I am changing the hardware I do not need to recalibrate my system or recalibrate my whole thing. So it is just one single shot of recalibration on the software and then your software is perfectly fine and working over there. So with that we come to an end of this practical application area discussion and one pointer which I would give to you is to look into this particular paper on retinal imaging and image analysis which had appeared in IEEE reviews in Biomedical Engineering and so this was around in the year 2010 if I remember.

So just have a careful look into this paper and there it is a very long paper where you have discussions on how retinal imaging is done the whole anatomy of retina. And then what are the different problems which are solved and more than a 100 papers which are discussed at length inside which will give you a very foundational overview, since we have a very short

span to discuss, I have just few pointers for you, but you can read a more on details in this particular paper as well. So with that we come to an end of this lecture and thank you.