Modern Computer Vision

Prof. A.N. Rajagopalan

Department of Electrical Engineering

IIT Madras

Lecture-39

So, what is normally done is you do we do something like you know like like the log right which is which is Laplacian of the Gaussian and and right which is which is of course, something that you do because it also gives you a control in terms of how much σ to use and then and then Laplacian of Gaussian you can even apply it at different scales. I mean I can take one σ which is very low and then I can see what is going on then I can you know increase my σ you can increase my σ. So what it means is that is you keep increasing the σ right the the structures that are very thin and all will start to will start to go away because because you already got to blur them out and only the very very coarse structures will remain right as you go higher and higher you know in terms of σ that is that is one way of you know one way of looking at things that means you have something like you know a control parameter that you can have. The other thing right that you still do is you still do not do not simply take all the zero crossings right what you normally do is you go and do a thresholding on the on the on the this one the magnitude right. So, because because you know the zero crossing right. So, so what you do is so you look at look at the magnitude of the gradient and then and then if you kind of do a thresholding ok.

So you kind of you kind of see. So, what this means is if there is something that is a very weak edge I think of think of strong edges and kind of weak edges right you always have strong and weak edges we sometimes call them as soft edges right. So, you may ideally want to get a capture all of them, but then there could be really something that is really weak right and then you know it is not even an edge ok. So, what you can do is you can look at the magnitude of the gradient and if it is and if it is right above above a certain quantity above a certain number then you can actually use that do a this one a thresholding and and then kind of leave yourselves with whatever else is remaining and that should give you a picture about what the what the most dominant edges are edges by mean by which I mean and you know edge this one pixel whenever I say edge point if you want to if you want to think about edge pixel or edge point.

Let me just show you some some images right on that. So that so that it becomes a little more clear and many of these things I have already mentioned ok you know during my during the other lectures and all, but I think we will just go through the ones that are that are relevant here this and all we have already seen it will not spend time this and all we have seen ok. So, this noise and all is here this and all this is very easy yeah right. So, here is where what I said is when you have σ right. So, you can see that you know the the fine edges and all will start to go away as you start you know right you know taking your σ up and this is an edge detector ok this is fine this is all I mean this is very straight forward, but let me come to so bell and all I have already told you ok now yeah right.

So here is where here is where right I mean. So, so one of the things right one of the things that that you can ask is can I not can I not do a thresholding of simply the this one right when I said that you can take the gradient magnitude. So, so let us say that something I take something like a so bell what is so bell - 1 - 2 - 1 1 2 1 right that is basically you know so bell kernel. So, suppose I suppose I took that I can find out to find out the this one I can find out g x g y I can I can find out the this one magnitude right. Now, what will happen if I simply if I simply did a thresholding of that.

So, let us say right I know let us say let us say I simply want to do that will that will that kind of give me the edges what will be the issue there if I tried that. Looks like that is also sensible thing to do because I am anyway saying right I want I want say thresholds I want magnitudes that are reasonably good those can be declared as edges what will be the problem there. One of the three things that I said will get hit when I said you need an edge right there are two three things that I said right an edge you need what is that what is that that that that might violate that will violate a single response. See, because what will happened is you know see when you when you when you when you plot the magnitude of of the gradient right what will happen is you know typically typically red you will have it is not like it is I mean so what will happen is red you will have something like this right if you kind of if think about the magnitude you have you have something like this right here is where the peak is now when you are when you are trying to apply a global threshold what will happen if you if you if you just make this magnitude too high red then many weak edges will go away if you make it too small then then you will be swarmed with lot of them. Whenever you adjust somewhere it what might happen is all of these there will be a whole band of guys which will emerge because you do not know where the where the local maxima is.

You do not know the local maxima and I just know that I have these gradients I do not know which one of them is a local maxima and all I just apply apply a threshold globally. What will happen if I make it too low then then then you will get more and more of these you know there will for an edge you need one response but then there will be collectively many guys coming out. Do you see that? So so so which is why because it is not easy to simply directly take the magnitude of of the gradient and use use only that information in order to do a this one you know kind of a thresholding because that single response thing that you really need which means that which means that you really ideally need right if you have an edge you ideally need a single pixel contour. You do not want a thick contour coming there right you need a single pixel contour of course you can have an edge rate that goes up and immediately comes down in which case you will have actually two edges. There is one that is going up and there is one that is coming down right so there you will have two they can be nearby that is ok but that is ok there you need actually two because there is one response for this edge there is another response for the other edge one is going up another is coming down you will have the theta and all accordingly right if you want if you are interested in that.

But when you have one edge you do not want you do not want a strip coming there right so so that is the reason why it is not so easy to simply take the magnitude of of the gradient and do it that is why you have got you have got something like there is an edge detector that is very famous that is called the canny edge detector right I mean it it works around this problem very smartly I will talk about it ok. Now now right so so what what ok now yeah unless you know the maxima and all right I mean this is still a problem therefore you can do a laplacian right and then in laplacian right see for example yeah so so right this is simply to show that it is going to be very sensitive noise and all that therefore you should actually blur it with a Gaussian ok this is laplacian of the Gaussian ok now you see right so what is happening is if you try to try to look at the zero crossings right when you when you increase your σ right you see that you see that your your zero crossings are kind of giving you more and more I mean I mean you are getting significant information right so so so when you when you traverse from here to here as you keep increasing your σ all the frivolous edges are going away because see if you just with a put everything out there right everything is active everything is important then then it loses meaning right so you really want to pick what are significant but you know what the problem is even if I showed you the right one the rightmost one which seems to have the highest σ in it it looks like it has picked all these dominant edges but even here if I told you it is a face would you think that it is a face suppose you did not I did not see anything else right now nobody showed me the person's face I just showed this image to you maybe some of you might say yeah I mean it is an image of a face but still it is not so clear right but then but then but then if I did this other one that I said if I looked at the magnitude of the gradient and I sort of you know applied a threshold and then kind of looked at it now you suddenly see a face right which looks much more cleaner now it looks like I am actually looking at a face compare this with the earlier one right again these are all subjective things okay that is why that is why right there is no clear way to tell which one one likes right somebody might like the other one perhaps right but then but that is the reason why right till this point we do not have something like you know there is a canny became became very sort of you know popular thing because it had some optimization inside here it is all like okay you do this you will get this but then there was a good motivation to actually write do it that way okay that is the way there is a way a log you know if you wanted to wanted to write implement a log based edge detector this is what you would do but then a more systematic way right in which in which an edge detection can be done is by is by do is by is by going through what is called a canny edge detector which I will talk about next canny is a person's name by the way John canny this is a pami paper I think it is 1986 or something okay. So this called the canny edge detector and this actually this is you know this is actually you know kind of a systematic way of looking at the problem and it tries to address that was all three things that means you know you should flag an edge only only only right if it has a significant kind of a gradient and the right it should have good localization and then it should have a single response right all three. So in a sense right so it actually solves it as some kind of an optimum and then it makes certain assumptions that the noise that is there in the image is actually AWGN very very simple assumptions and then and then right we are not we are not kind of afraid you know we are not we do not have time to go to the details of how he arrives at it but then eventually shows that the best filter to actually to actually use under that condition and then the optimization it is in terms of you know what you call this one this one a product of what you call signal to noise ratio versus the versus the localization okay and that that product he says it you know it becomes maximized when when you use a particular kind of filter what do you think might that might that filter be that he arrives at as the optimum filter for doing edge detection right you have seen so many filters now which do you think which you suspect might have been the one that he that he came up with a derivative of the Gaussian okay dog not the difference of Gaussian a derivative of Gaussian okay. So he actually analytically shows this okay now this is more formal okay this one this work is much more formal than the ones right that we that we saw till now so this is a derivative of the Gaussian but but then right this by itself okay is not really really enough I mean he has some he has some nice ideas okay you know right in that work so derivative of the of the Gaussian and that makes sense right because he has assumed noise and therefore it that has to be a countered so we need some kind of a smoothing and he believes that you know the best filter to actually use that and is this and then the signal to noise ratio also right is it is kind of you know a different interpretation it is about the it is about the same number of the positive edges right that you actually capture divided by divided by the by the by the right number of false edges that you flag.

So signal to noise ratio is not our usual sort of signal to noise ratio he kind of constructs the problem around around around this notion of doing an you know this an edge detection and so this is the right three concerns are localization right strength so so that means right it should be it should be what the I mean it should be a proper edge and then the third one is single response right all this has to happen okay. Now these two seem to be okay but then the strength right by which what he means is that even if you have weak edges and so on it but if they are relevant they are they they should be picked up okay that is okay right that is what that is what he means. So so what is relevant and what is not relevant right we will see now okay so the main steps of canny of canny edge detector and this you will find in your matlab and all right this is a very open CV everybody will have this smooth the image with with a Gaussian filter then compute the gradient compute the gradient magnitude and direction. Now okay now what are the things right that that I think I forgot to mention is that in a Laplacian it is an isotropic operator it does not give you a direction information right you saw that right you do not have a direction information you just have the zero crossing information there and we and we get a views that right so but then here at compute the gradient magnitude and and it is and it is you see direction right. So now finally right we are again coming back to this that that is the theta that I mentioned right in the beginning so that this guy actually actually uses it okay compute the gradient so this is like your dog smooth with a kind of Gaussian and then take a difference and then compute the magnitude right and this one okay direction using a discrete approximation and all that that we know right how to do that.

Then apply apply non maxima suppression has you guys heard about this this NMS in some context this is not this is not something particular to edges or something any other context has any of you heard about non maxima separation perhaps right you must have okay so so kind of non maxima suppression means that when you are kind of looking at several guys coming up so so you want to you want to really throw away the guys which are really not a maxima so you suppress them. So so you can see that that single response thing right has to kind of right do with this because if there are several things that are that are trying to trying to trying to kind of right trying to tell that we are all important you want to say the only guy who is really the maximum should be important rest of the guys should should just go away right so that is why this maxima of suppression and this NMS right will again revisit later in a kind of a different context and so on to the I mean gradient magnitude. Any gradient okay no any gradient value okay yeah right this is like a this is like a succinct summary of the canny edge detector any gradient value that is not a local peak along the direction of the gradient again this is again important along the direction of the gradient is set to 0. So then what does it mean so for example, so what this means is if I had an edge like this right which is kind of going up from from a kind of right right low to high and if you were to if you were to compute the magnitude right a gradient you will find that you know right within this region there will be there are several guys which will have a significant value if you just directly did a thresholding many of them right may still actually survive. So what it is saying is saying is along the along the say direction of the increase right which is what which is what is the orientation telling you along there you try to find out right which is which is which is which is supposedly a true maximum and then right rest of them should go away.

So the way right so the way it is done is suppose I take a take a grid like this where I am and suppose suppose right I am suppose I am I am I am at this at this pixel location right and I want to know if the gradient magnitude here should that be suppressed or not suppress means set it to 0 that is right it just goes away. Now in order to do that right so what you have to see is what is the what is the actual this one right this one edge orientation there suppose the edge orientation is like this at that at that pixel it could be anywhere right whatever it could also be in some other direction some direction if you choose then kind of what it is saying is right if I were to compare this with this if the if the orientation is that if the edge orientation is that then then kind of right what it is saying is compare the magnitude of the gradient at let us say suppose I call this some a b c right suppose I come suppose this is a this is b this is c then what it is saying is if the if the right gradient of the magnitude at at a is greater than both c and b then only then only right we will actually leave it there if it happens to be one of the others is greater than that then this goes to 0 a goes to 0 we are only focusing on a we are comparing the magnitude of the gradient at a with b and c it survives only if it is actually right you know greater than both c and b. Now of course it can also happen that that that right I mean that for example it can also happen that the edge orientation is such that you end up end up something somewhere like that because we do not know which way the edge is oriented so the edge is oriented that way then right you might actually you might actually end up in a zone where you do not have a direct value I mean if you if you if you are lucky you could have a situation like that right where you actually end up in a in a sort of an integer pixel grid you need not right and then what you would do you do the standard thing which is again right doing some form of some form of right interpolation you can do a simple linear interpolation use these two values whatever is here whatever is here and then use how far away is this location just like I did know the earlier case something similar to that assume that right it is kind of a delta away from the top pixel 1 - delta away from the bottom pixel do a linear interpolation get some value here and similarly get some value here again compare this guy with those two suppress under those conditions otherwise otherwise it stays on right this is this is what is called non maxima maxima suppression and you can see that it is making use of the gradient magnitude as well as the right orientation okay both both are both are required but then this it is not over yet the the other key step right in canny is this one after you have applied no no okay there is something called a thresholding which you expect will happen okay because all the while that we were kind of talking about a thresholding and linking and this is typically called a kind of a hysteresis operation hysteresis you all know okay hysteresis is something kind of you know a memory right something that you remember and so on so here right there is some kind of a you know a memory involved and we will see what that is and I like this step right in the sense that this is one of those one of the steps right where you sort of make sure that you do not lose weak edges okay so this so the steps are like follows right so what it does is mark mark as edge pixels okay because right we have still not said anything okay we only suppressed the things that are not a local that are not non that are not a maxima see those have gone it is not like all of them are edge pixels okay now only we are coming to what we would like to declare as an edge pixel mark as edge pixels right okay I think this is the first step mark as edge pixels all pixels with M i j okay which is what is M the magnitude of the gradient at i comma j greater than some this one threshold t 2 now this is where this is like a hyper parameter okay this you have to you have to select okay you have to or maybe well you can do you can you can there is no there is a straightforward way right to to arrive at it because depending upon how you choose t 2 you might actually end up with with either right a little more edges than what you want or a little fewer than what you want okay this is a hyper parameter t 2 then mark as non edge pixels that means these are these are really not edge pixels non edge pixels all pixels with M i j less than t 1 some other so so this is like a high threshold this is like a lower threshold which means is that something that is weaker than that don't even don't even don't even accept it throw it away and some people say t 1 can be taken as t 2 by 2 but again these are all heuristics okay some some other heuristics of course is involved then what about what about the say rest right you have you have these guys that are lying between between say right t 1 and this one t 2 okay we are still undecided about them what would you do those are the those are the weak those are the soft edges right because you have something like a very clear edge which you can which you can clearly see you know exceeds a certain threshold clearly declare as edges something which is very weak very clear goes away right what do you do with this with these right middle guys you want to be able to use them you don't want to throw them away so so so actually it's smarter than that i will write it down okay so mark as edge pixels all pixels okay mark again mark as edge pixels all pixels with m of i j greater than t 1 but less than t 2 that are 8 connected 8 connected to at least 1 pixel with m i j greater than t 2 along a chain of pixels this this looks a little verbose but i think chain of pixels with what do you think with m i j greater than t 1 what does this mean see when you say 8 connected right so what it means is when you when you take you know a 3 cross 3 window when you are sitting here there are there are 8 of these guys right around you that's the that's the 8 connected neighborhood right if you take a simple you know 3 cross 3 window so so what it is saying is right between between so so all those edge all those pixels that have a value between t 1 and which have a magnitude of the gradient that's between t 1 and t 2 right so what it is assuming is that in the real world right there is going to be linking up right there will be like you know there is going to be a continuity right and so on so what it is trying to say is if a path exists from that pixel right which has which has let's say a gradient between t 1 and t 2 but if a path exists which should be which should be right within within this 8 pixel neighborhood and and and right and through this 8 pixel neighborhood if you can arrive at some other strong pixel but then when you when you go through this right i mean you should be able to you should be able to go through go through pixels that have m i j m i j greater than t 1 it should not be through the through the other guys which you have already knocked off so it means that if there is a weak link right because this because the t 1 and all like like you know just about t 1 are all kind of right weak guys right so what we are saying is if if a weak guy has a link to a strong guy through a chain which can which can be which can be a simple local neighborhood why why local because edges and all will have some will have some spatial relationship right so what it is saying is if that path exists flag it because very likely that the likely you threw you threw it away because there was a gradient that was weak or something and then you seem to think that it is probably not an edge pixel but in reality it might be there because that is how natural images are formed right there is always you know a continuity and so on so so so this so this idea of thresholding and linking so linking is a step where you link a pixel you really declare it if this happens otherwise you say you say right it is not an edge pixel i mean right so actually clean idea right i mean make sense ok and if you do like this right and you should ok i will just show you a few examples right just to show how canny works i hope i have them yeah this and all is there so derivative dog of so here you have dog non maxima suppression this and all i told you already so right so so here is how it is so if you do a non maxima suppression see this see this kind of you know a thickness that you see there right that is what i meant by not a not a single response right and and see this right after you do this non maxima suppression right you kind of see removed all this all those thick guys and then something should happen is that the last one no ok then then you then you go through the strong edge weak edge ok this one for some reason then then you then you do this ok and effect of effect of the gaussian so again the same thing right that we had there we can again have oh i think this is not in full screen ok so again it what we had there right we can also have here so for example right you can have a σ that is small σ that is high again you can run canny at various σs because there is a derivative of gaussian sitting there right and you can say you can see that the edges that come out are actually are actually quite nice and and you know you can have situations where you want to combine the edges and so on right we are not we are not going to get up to get into get into that but as you can see right it is a more systematic way of kind of using all that we learnt in terms of gradients orientation and all that other than brought together in a sort of a nice framework in order to solve a fundamental problem which is this which is which is which is one of finding edges ok I will stop here.