Modern Computer Vision

Prof. A.N. Rajagopalan

Department of Electrical Engineering

IIT Madras

Lecture-46

So, we will go towards shift, but then before that, okay, there is something that I would like to talk about what is called blob detection, as I mentioned in the last class. But it is not just, you kind of look for corners and images have various things and you want to look for some things that could be generally of interest and one such thing is a blob. Blob is like any area of where let us say there is some activity and there is some interesting stuff kind of going on, right. So just to kind of give you an idea, right, as to and why for example, why we need this, you know, why we need to be able to handle scale and all that, right, will be apparent now. Oh, before that, right, let me just show you a few examples of that corner detection. I think last time I forgot to do that.

Let us just go back to last time was what, 21, right. So there, I think there is something that I forgot to show at the end. So we went up to this and yeah, so this is something that just so that it stays in your mind. So you have kind of two images, right and you want to be able to match key points on them in this case corners and what you can do is you know, you can compute a corner response and of course, this by itself is not useful.

Then you can do find points with a corner response greater than a threshold which is what we said last time, right, so that the weak guys sort of drop off. Then we have to do something like a non-maxima suppression which is like you know, pick only the local points of local maxima and I hope you are able to see all these small points right which have emerged, those are all the corners that it has picked and then if you actually plot them on the image, right, so kind of they kind of look like that which is also which is what I meant. I mean a corner the way we define right is like what we think right it should be like, but eventually anything right that would actually give out a corner response that is sufficiently strong right it will pick all of that. So it is for us to you know, it is not humanely possible for us to sort of examine it you know what kind of a corner it is and so on, but then wherever the corner response has been significant right it has picked up all such points and of course, you know one of the things that I also mentioned was the fact that this rotation invariance we already saw right when you can have a rotated picture and then the features are actually covariant, but and of course, you know you should be able to also match them. A scaling rate as I mentioned not invariant scaling and anything above that and you can imagine why so because I mean if I keep zooming in right I mean this is maybe at some scale of zoom and then if you zoom further right then maybe you know it will become even more smooth and then maybe that kind of a corner response that you would get I mean you may still think that there is something like a bend there would that not be a corner or something, but then the strength there would automatically go down then it would not qualify for a corner with sufficient strength right and so which is why you know which is why we say that it is not actually covariant because you know scale depends upon at what scale you are examining it and there could come a point where something that was originally dictated as a corner at some scale no longer looks like a corner and therefore, we cannot say that we can take any scale and then right and then be able to say that it is a corner.

Okay, so I think this was just a there is something that I wanted to show before we actually proceed to blobs. So, let us come to the blob thing. So, corners of course you have already seen so that is so let us go and see what is the so blobs right are okay so here are what we would consider as blobs. So, right things which have which have a lot of activity under them and of course you know this is for some threshold and so on. So I mean wherever it you know it picks up so as you can think of so you can think of right some kind of a template right some kind of a structure that you have which you should be able to apply.

But then one of the things that you actually notice is that you know there are these blobs at various different scales here right it is not like there is just you know one scale of one scale of activity the activity is going on in different different scales. And you need something like a probe you know which will kind of probe them at different scales and try to tell at what scale right something is actually still something can still be considered as a blob of activity. And maybe write some other scale it does not look like it is so strong but then at some scale it appears to be a blob of activity. For example, right I mean if I take if I take this guy right here okay now if you see that circle right that has a certain scale right. Now if I had taken a circle of this kind right which is much smaller okay and it is quite possible that you know that at that scale right it did not flag it to be a blob right I mean in the sense that maybe it had some strength whether it still did not consider it to be high enough.

And then right as it worked towards a different scale at some scale right it sort of flagged it and it said well at this scale right it looks like a good fit and then right that is like a blob. And then again right so again maybe if I cross that scale suppose I take something like this and then and then right if I were to put that there right for some reason it does not think that right that is a blob right. So at some scale it kind of so it is like a probe you know so where the probe should sort of you know automatically tell that you know when it actually probes it then for a smaller one it should probe at a probe with a kind of a smaller. So it is like you know what to say right. If you have a big probe right and suppose you want to get a say resolve 2 points you cannot right I mean as the area of the probe increases your ability to be able to say that there are actually 2 points that are separated from each other will go down right.

Somewhat similar to that so when you want a probe then the probe size should automatically adjust so that when you actually put it then for a smaller one the probe should automatically become small and then say that yes right this is a blob. If it is a big one then the probe should automatically become big and then sort of say that well at that scale right this looks like a blob and so on right. Yeah I mean right that is a kind of you know analogy you can have because I was thinking whether I can give you something from you know Fourier this one but I do not think it is directly the same. Okay we will anyway we will see that so blobs right. So okay hey wait a minute this is supposed to have some animation.

No I do not see that okay anyway right so what happens is actually there is an animation here okay or maybe this plot does not have maybe okay so it is like this right so at this scale right so if you see the x axis right that is a scale and you have seen right that is actually taken one is kind of zoomed in another is zoomed out right this is like from far and then that is like a little closer. But in both cases we know that right it is some object of interest sitting there and you see the red at some scale it seems to flag this so because of the fact that this object looks bigger therefore right it requires a larger probe and therefore right at some higher scale it actually flags it. It does not mean that at other scales there is no flagging it just that right at some point it kind of hits a peak and that is where you get the maximum strength and similarly right here for example because of the fact that this is actually zoomed out so it is kind of smaller right and therefore you know and what do you call right and therefore it follows that the scale that you require is going to be probably much smaller and therefore right at a smaller scale it flags it whereas it is not for example if that scale of 10 right wherever it is I mean maybe somewhere here it you see that it is still giving out a response but that is not the scale at which the object is able to give out the maximum strength right with respect to the probe. So the whole idea behind right you know behind doing shift right which is what we eventually want to do is to be able to handle so the transform itself is called scale invariant feature transform and of course and it should also have you know other characteristics right I mean the way it is been built is it is not something built just for scale even though you know by the name of it it looks like something that can only deal with scale but it can do several things okay and it also depends upon how it is evolved and then how they come up with actually a descriptor that can handle various variations and so on but the blob detection is like the first step right to even understand why shift works the way it does I mean for example shift uses you know a difference of Gaussian and all that is a paper by a person called Loewe it was a 2004 paper and apparently there is some interesting anecdote to that maybe some other day but then he is the one who actually based upon this blob and all right so he kind of advanced it right and that is how shift came about. So for example right so the basic idea okay now yeah I think all right I think you know before we see these pictures I think let us let us kind of do some do some analysis just to understand it what we need okay.

So blob this one detection yeah so I think did I also show the flower example I did right okay now in order to understand it let us kind of look at a 1D illustration because a 1D illustration is much more easy to easy to understand and it is easy to see what an edge is and so on it is so I am going to take 1D as an example but all that I say can easily be you know can easily be extended and analyzed in actually 2D okay. Now in all of this right if you see they will use what is called a normalized Laplacian of the Gaussian okay you have seen log right there is something called a normalized Laplacian of the Gaussian okay. Now this normalization right there is a certain factor by which you multiply in order to be able to normalize and one has to understand as to why they do it okay and why that normalization factor comes in and so on. So maybe what I will do is you know I will again go back go back to the to this to this example let me see where is that let me go to this example okay. Now if you see right here is an edge okay has some noise and so on not really an ideal edge right has a little bit of slope and then goes right not it does not have infinite slope as a finite slope and then hits a higher value.

Now we know that right we can actually probe it using this one a Laplacian of log right so we know that if we can identify the zero crossing then we know where the edge lies and therefore right if we had you know if you were to take a what is this G is actually a Gaussian okay. So this is like a second derivative of the Gaussian which is what we mean by a Laplacian. Now in 1D right this is how it looks like so this is what we call a right Mexican hat okay so in 2D right you can imagine right it looks like a Mexican hat right so in 1D that is how it will look. Now if you were to see convolve it right with this signal F with the step then we know that right this is how this is the kind of right when a response you will see I mean right this we know this we have also plotted sometimes right and we know that the zero crossing is where the kind of right edges. But now right if you go to the situation right where I have a blob see a blob right in 1D right if you were to take this one a cross section it will be like you know some intensity and then it suddenly goes down to 0 then again comes back up right.

So if you were to kind of think of a blob right then really it is kind of you know a double edged oops double edged this one right it is like a step right I mean you could have the other way around also and so here it is like going from 0 to 1 it does not really matter. Now if you try to probe this okay now right what we plan to do is we will take a log right and we will try to see right I mean how it actually responds to these edges. Now at this point right I mean you know you see that right a response right which you will get for that edge is this and a response right which you will get for this edge right is going to be this and of course you see that the signs are opposite because of the fact that in one case it is a rising edge another case the falling edge. And this is the kind of zero crossing right is somewhere here okay that is where the edges kind of say located. Now this by itself right it seems to say something but then right you know I am still not able to able to tell much.

Now if this pulse or this width of this blob right if it shrinks now if I make it smaller what will happen then these guys right will come kind of right will come get to say closer now. Then again right I kind of reduce the width further then they come even closer and then right at this point okay when I kind of when I kind of bring them even more close right I see something like an say extremum now right it could be a minimum it could be a maximum but there is an extremum there okay. Let me see there is another slide now this is how this is what is an ideal situation okay which wherein let us say right at some scale things have matched and at some scale right think about think about the probe right matching with the width of this blob and then flagging it right. So, that at other places there was some strength but then but then there are but then for this probe the actual signal that it can actually pick is this blob you can also interchange the roles right you can actually think of think of right I mean you can actually think of think of having the blob to be of a fixed size and then you can actually probe it with the log wherein the sigma can be then varied and you can then ask at what sigma rate does it match the blob right. So, this can be interchanged now actually you know in reality what happens is this right you have a blob right and you are and you are trying to trying to write you know this an examine it now when you actually now in this case right in this case we have actually interchange the role right we have kept the blob size to be fixed and we are trying to change the state the sigma of the log right but that is why log is actually good I mean it gives you a control about what you want to do right.

So, there is a sigma that allows you to control now when you when you actually increase this sigma rate what you find is that is that right is that of course, you know is that is that right you go from here then you go from here to here then you go from here to here and you actually see something like this one a minimum there and then and then it goes from there and then again right it seems to rise after that ok. It is not very clear from the plot but then it is not actually then it looks like something is going on there at some it looks like right at this point it probably clicked the best right in terms of achieving something, but then if you look at the strength the magnitude of or the or the strength of the response that seems to have gotten scaled down right as I kept right increasing my sigma. So, in a way so that is called unnormalized Laplacian right that means I have not applied any kind of normalization I have just applied a Gauss the log and then just kept on I kept on increasing my sigma and what has happened is it is of course, you know brought these things close together they start interacting at some point and then at some point it looks like something nice happened here, but then right I cannot say that you know I have already achieved an extremum because the way the range from here to here is not the same right this is been you know a difference and after this something else seems to be happening maybe if I increase it further you know maybe right this will become even more flat and so on. The ideally right what you would ok. So, so ideally right what you would like to have a some kind of an you see some sort of an you know what you call you need an you see invariance.

So, that the strength is not actually affected by my sigma I mean that is when right we will be able to ideally pick up pick up as to when when the probe actually matches the blob right in terms of the scale ok. This is what this is what right we would like to see in X as to how that is what is actually a actually a normalized Laplacian of the Gaussian ok that is the role of the that is a that is the reason why we go for a normalized log because an unnormalized log is not really the best thing to use.