Introduction to Robotics
Professor. Balaraman Ravindran
Indian Institute of Technology, Madras
Department of Computer Science
Lecture 7.2
Recursive State Estimation: Bayes Filter

(Refer Slide Time: 00:16)

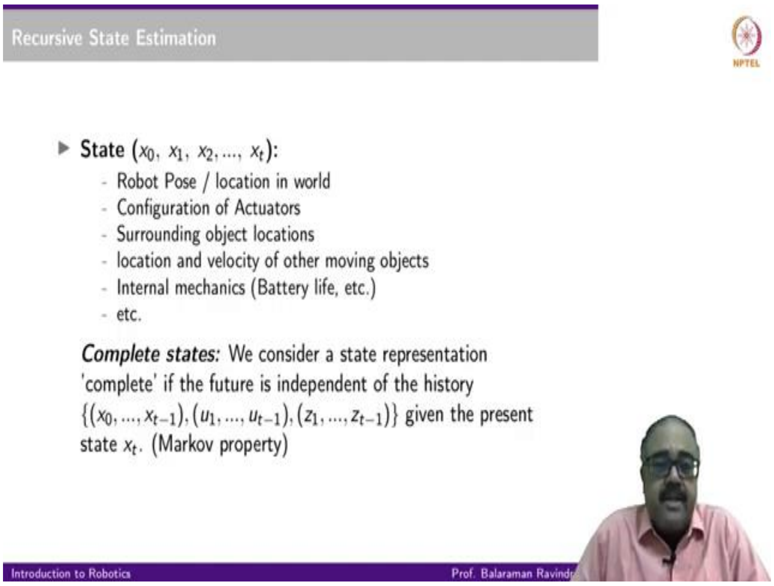

Last lecture, we were looking at state estimation, we started talking about Recursive State Estimation. And so, we were talking about what constitutes a state. The state at each time t, could be a very complicated vector of various entities that you could record, like the robot pose, the location of the world. This could be the xyz location, and also the orientation theta.

And the velocity with which, if it is a mobile robot, the velocity with which it is moving, and so on, so forth. And then the configuration of actuators, so in what angles are the arms in the gripper is in, and whether it is holding an object, and so on, so forth, and locations of surrounding objects, etc, etc. So, we said that this could be potentially a very complex state.

(Refer Slide Time: 01:05)

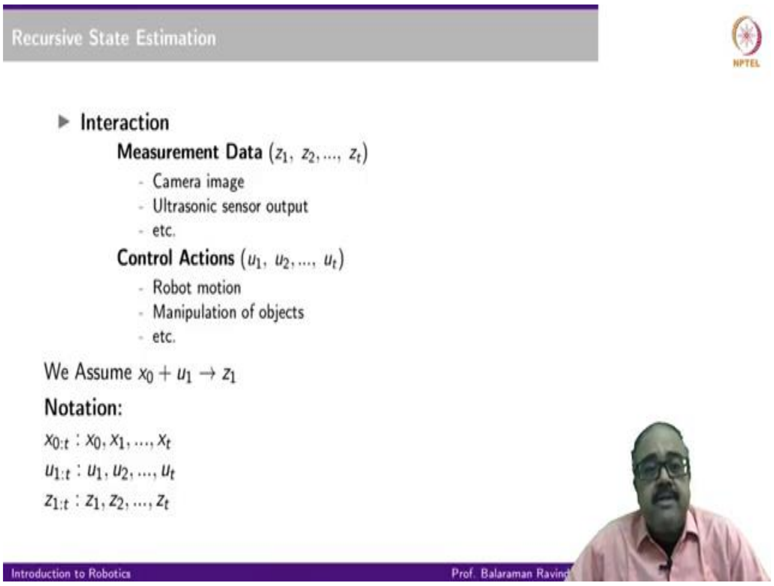

And then we started talking about measurement data. This is essentially what the sensors give you and likewise, we had z1 to zt. So, this measurement data could be camera images, could be ultrasound sensor data, and could be other kinds of touch sensors and internal indicators, like battery levels, and so on, so forth. And then finally, you have a set of control actions, which could be movement actions, it could be manipulation actions, and sometimes it could just be sensory actions, like turn on a camera or rotate a camera in a certain direction, so that you get a different kind of input.

And so, we also adopted this notation that you start in state x0, you perform action u1, and you end up at x1 that were you take measurement z1. And so likewise, throughout, you are at xt minus 1, you do ut, you end up at xt when you make a measurement zt. So that is the way we are going to be looking at it. And the notations are as follows. So, when I want to denote a entire sequence x0 to xt, I will use x 0 colon t, just remember that. Likewise, u1 colon t and z 1 colon t, to denote the entire sequence, it could start from anywhere and could go to anywhere. It could start from t minus 1 and go to t plus 1 also. So that does not matter.

(Refer Slide Time: 02:34)

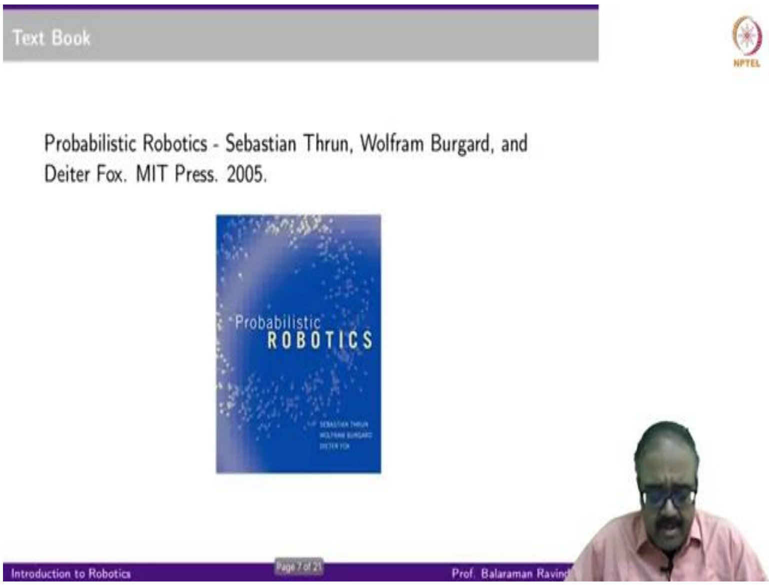

So, in the last lecture, I mentioned that we will be using this textbook, I just wanted to make sure that everyone has gotten down. So, we will be using the textbook on Probabilistic Robotics, by Sebastian Thrun from Burgard and Dieter Fox, it is from MIT Press and there are draft versions of the textbooks also available freely online. This is not the complete version of the book, but as a reference, you could use these draft versions. And the book is very extensive and like I said, we will not be covering all parts of the book, only some of the highlights from the book.

(Refer Slide Time: 03:10)

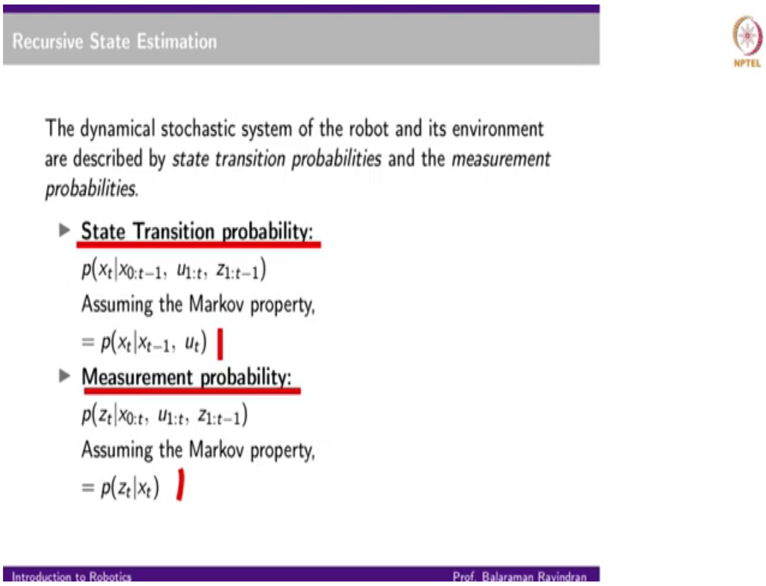

And what we then started talking about is the system model. And I said the system model consists of two quantities, one is the state transition probability, the first one is the state transition probability. The first one, so the it consists of two quantities, the first one is the state transition probability, where you look at the probability that you land up in a certain state xt, for example, probability that I land up in front of open door, given that I have been in the past sequence of states given by x0 to xt minus 1 and that I have taken the actions given by u1 to ut and I have made observations z1 to t minus 1.

And then we assume that the Markov property holds and therefore we can write this as p of xt given xt minus 1 and ut. And the next quantity we look at is what we call the measurement probability, which is assuming that under the Markov property, the measurement probability is just assuming that I have, I am in state xt, what is the probability that I will make a measurement zt. So, we are looking at probability that I will be making a certain measurement zt at time t, given the history of states that I have visited up till time t.

And the actions have taken up to time t and the observations I have made until time t minus 1, and all of this put together or the factors that my current observation could depend on. And if you are assuming the Markov property, the observation depends only on the current state xt. And we also talked about how zt does not figure in the state transition probability, because zt does not cause xt, zt is caused by xt that is actually captured in the measurement probability.

Now, assuming that you did not have this noise, assume that the world is clearly observable. Assume that all you need to know, let us say you have a wheeled robot that is moving around in a 2d workspace, all you need to know is the exact x and y coordinate of the robot. And that is all the information that you need to make all the decisions you need in the world. In such a scenario, let us say, if I make a measurement, I know exactly where I am. So, I know my state because the measurement is going to be like something like a GPS signal, that gives me the xy lat long, very accurately.

So, at every point, I will know what state I am in. So, if I tell you that you have made a measurement that tells you your lat long is x and y, then I know that I am in location x, y, then I make a movement, I say I move north and then I make another measurement. Now this measurement again gives me my lat long x and y, and therefore, I know exactly where I am. But the entire complication, the whole reason that we are looking at recursive state estimation right now is that I do not have such noise free measurement and I also have a problem with my modeling, motion modeling.

But right now, we are only looking at the fact that my measurements are not noise free, and they are noisy. And therefore, I cannot exactly tell you what state I am in, the robot is not able to decide what state the robot is in, and therefore has to, so the robot is not able to exactly measure the state it is in, and therefore has to look at a probabilistic estimate.

(Refer Slide Time: 07:05)

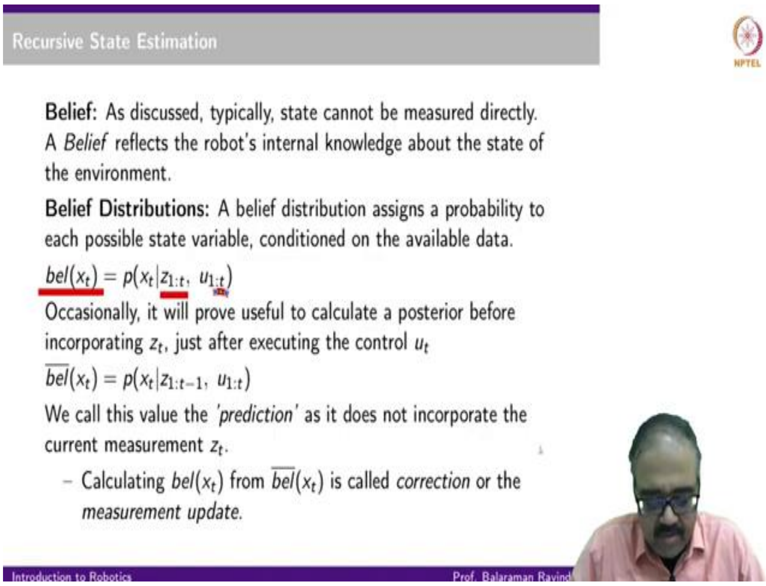

So, we call this estimate of the robot, the robot is essentially trying to, you know decide what state it is in. At any point of time, this internal confusion about the actual state the robot is in, is captured in what we call as a belief. The belief is essentially reflects the robots internal knowledge or internal confusion, if you would be the flip side of it, is about the state in which it is in. So, for example, if the robot could be in one of two places, it could be in room 1, or it could be out of room 1, let us say these are the two things the robot could be.

And then when it makes a measurement and the measurement just tells it, you are near a door, that does not really let you know, whether you are inside room 1, or whether you are outside room 1, it could be near the door in any way. And then the only way that you could be sure about where you are is actually take some actions and see what happens with regard to the effects of actions, and then continuously refine your belief. So, when I start off with no knowledge of where I am, I could say that there is a half chance that you are inside the room and a half chance you are outside the room.

So, this kind of a probabilistic distribution over possible states is called as a belief. And sometimes we use belief, sometimes we use belief distribution. So formally, a belief distribution assigns a probability to each possible state variable and conditioned on everything that you know so far. So formerly, I would say that belief xt, so belief xt would be, belief xt would be the belief that you have that, what is the state that you are in at time t. So, xt is remember, xt is a variable that tells you what state you are in time t.

So, it could be in multiple states. So, it is actually given by a probability of you occupying a specific state xt, given the entire history of observations that you have made, and the entire history of actions that you have taken. So, remember that you never get access to your actual state. Sometimes you can assume that you know x naught, that is the state that you start in, but even that is not available to you often. So, you basically only have a sequence of actions that you took, and the sequence of observations you made, that is all that the robot knows for sure.

And so, given the sequence of actions and observations, what is the probability that I am in a particular state at time t, so that is essentially the belief at t. So, for every possible value that xt can take, you will have one probability, and that basically is your probability distribution at time x, at time t. So, this essentially, is a quantity that we call as a belief, and we will denote it by bel x t. I think I know that it is a little confusing right now, it will become clearer when we look at an example later.

And just keep in mind that when I say belief xt, it does not really mean that the robot is actually in two different states with some probability. The robot is always in one state, this robot is not a quantum robot, that it can be in multiple states at the same time, robot is always in one state, just that it does not know what the state is. So, the belief does not encode the actual position of the robot, the belief represents the robot's estimate on the position.

So, believe tells you what the robot thinks is the position, not what the actual position of the robot is. The actual position of the robot is some x in the world that you do not know yet. So and the belief tells you, what is that robot's estimate of where it is in the state space. So, while we are trying to compute the belief, it might sometimes be useful for us to compute a quantity, which we will denote as bel bar, which is a prediction of where it expects to be after it does an action.

So, the bel bar xt is essentially the probability of xt, given z1 to t minus 1, and u1 to ut. So, the difference between bel and bel bar is that bel is conditioned on z1 to t and u1 to t, where bel bar is conditioned on z1 to t minus 1 and u1 to t, it needs to know the last action that you performed also. So now you might actually start thinking a little bit about, hey what happened to all this Markov properties week are talking about, why are we talking about the belief, now conditioning it on the entire history of observation, the entire history of states in the entire history of actions, and so on, so forth.

The reason that we have to condition on the history of observations and history of actions is because we do not know the state. If I know xt minus 1, then I can make this more Markov. If you know xt minus 1, then it depends only on xt minus 1 and ut. Since I do not know xt minus 1, I have to look at the entire history of observations and the entire history of actions I have performed in order to define my belief.

And note again, the dynamics, the underlying dynamics of the robot problem is still Markov, we are not changing that, the underlying dynamic system Markov but because I do not have access to the true state, my belief estimations would be dependent on the complete history. So, you should remember that. So, just because I put the whole history here does not mean that the problem has become non Markov, is it clear. So now, so we have this quantity bel bar, which say is the prediction.

So why I say it is a prediction, I have not yet made a measurement of where I actually landed up. So, I have data for all the past measurements I have made, z1 to zt minus 1. And I have data of all the actions I have taken, including the current action. So, after I have taken the current action, I am going to say, hey, where should I go, I should be, you know, maybe I am trying to leave the room. So, with some probability, you know 80 percent I will leave the room. So, I would say that, hey, if I was in the room earlier, now I have left the room. And that is what I should be looking at. So, I am no longer inside the room.

And this prediction, as you see on the slide does not incorporate the current measurements zt. So, if you notice the difference between bel bar and bel was that measurement zt. So, going from bel bar to bel, so calculating bel from bel bar is actually called the measurement update, or sometimes called the correction update. So, whenever I go from my, whenever I compute bel bar, I am saying that I make a prediction, so we will we will make this more clear in the next few slides. Whenever I compute bel bar, I am making a correction, I am making a prediction.

And when I compute bel from bel bar, I am making a measurement update or a correction update. So sometimes bel bar is called the prediction or a movement update or motion update, or the transition update. And going from bel bar to bel is the measurement update. So, we saw these two quantities earlier.

(Refer Slide Time: 15:24)

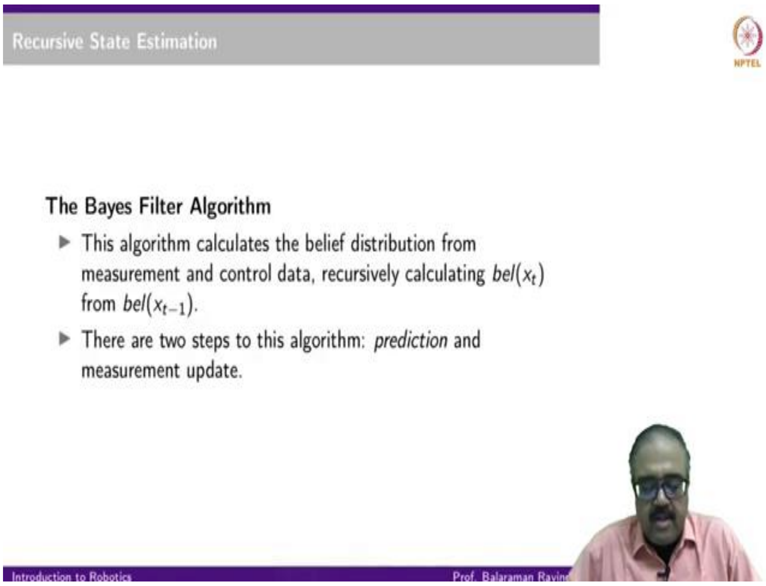

So let us, let us move on, so let us move on. So, we will now look at the base filter algorithm. So, we will now look at the base filter algorithm. And the idea behind the base filter algorithm is to recursively update your belief. So, the idea here is that I am going to compute belief of xt, that is, that is the belief at time t. So, what is my distribution over the value that the state can take at time t, I am going to compute that from not only the dynamics, I know the motion model, and I know the measurement model.

We assume that I have the motion model, I have the measurement model and I also have access to the observations and the actions. So, what do I have, I have the motion model that we talked about, the transition model that we talked about, I have the measurement model, I have access to all the observations I have made so far, and I have access to all the actions I have taken so far. Given that I have all of these, how do I take belief xt minus 1. So, what is belief xt minus 1 the distribution of states at time t minus 1.

How can I take that and use that to compute belief xt. So, and as you could guess there are two steps to this algorithm. So, we have the prediction step where we compute bel bar of xt from bel xt minus 1, and the measurement step that we compute bel, xt from bel bar xt, is that clear, so let us move on.

(Refer Slide Time: 17:03)

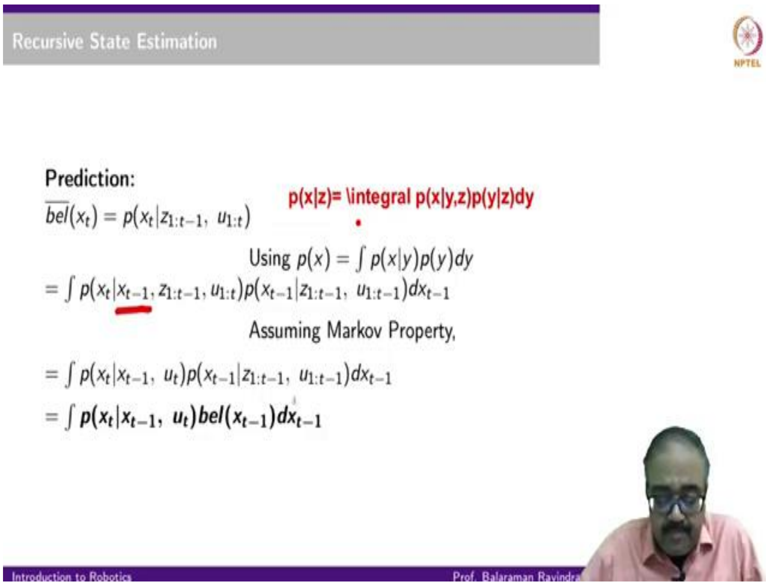

So, let us look a little bit more detail at the prediction problem. So, what is a prediction problem. Remember, a prediction problem is compute bel bar xt as a probability of xt, given z1 to t minus 1, and u1 to t. Now, we know that from identities of probability, we know that p of x is equal to integral p of x given y, p y d y. Now, I am going to use this in order to simplify this expression, and somehow introduce bel xt minus 1. So, to simplify this expression, so I am going to do this integral over all possible values that xt minus 1 would take.

Just to recall, so what this integration does, it is essentially runs over the entire gamut of values that y can take and look at for every value that y can take, what is the distribution over x. And then multiply that with the probability that y can take that value, and do this for the entire range of values that y can take. So, this integral would give, allow me to condition x on y and give me what is p of x. So likewise, now, I going to take this expression, so I want p of x given z and u, so I am going to condition it on xt minus 1.

So now we have a probability of x given y this integral, property x even y into property y. But suppose I want to do probability of ex given z. Now, I can write that as integral of probability of x given y comma z times probability of y given z dy. And so of course, I had to take the integral over here, I can take the integral of that, and that gives me the expression that I want. So that is exactly what we are doing here now. And once now that I want probability of x given z comma u, I am going to do probability of xt given xt minus 1, z1 to t minus 1 u1 to t times probability of xt minus 1 given z1 to t minus 1, u1 to t minus 1 dx t.

So, if I assume that we actually have the Markov property, so some of these things can get simplified. So how do I simplify with the Markov property, notice that I said the Markov property applies whenever I know xt minus 1, if I do not know xt minus 1 I have to condition on the entire history. But if I know xt minus 1, I can get away with throwing away the history. So, since I have xt minus 1 here, I can throw away the history.

And then I can simplify that to probability of xt given xt minus 1. I still need ut because ut comes after xt minus 1. So, I need xt minus 1 and ut, that gives me this the first expression using the Markov property. The second expression stays the same, it is probability of xt minus 1 given z1 to t minus 1 u1 to t minus 1, dx t minus 1. But then if you think about it, what is this expression probability of xt minus 1 given z1 to t minus 1 and u1 to t minus 1.

(Refer Slide Time: 20:52)

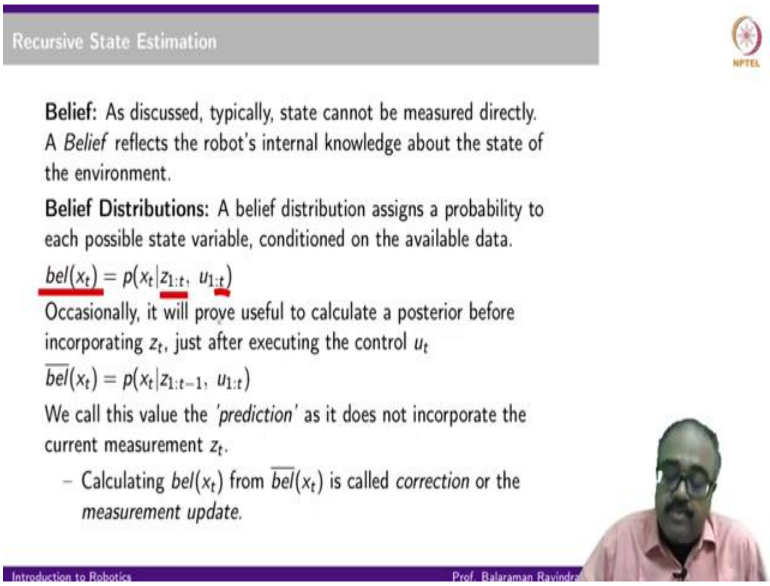

If you think about it, that is exactly what our belief expression was, belief xt is probability of xt given z1 to t, u1 to t. Now belief xt minus 1 would be probability of xt minus 1 given z1 to t minus 1, u1 to t minus 1.

(Refer Slide Time: 21:08)

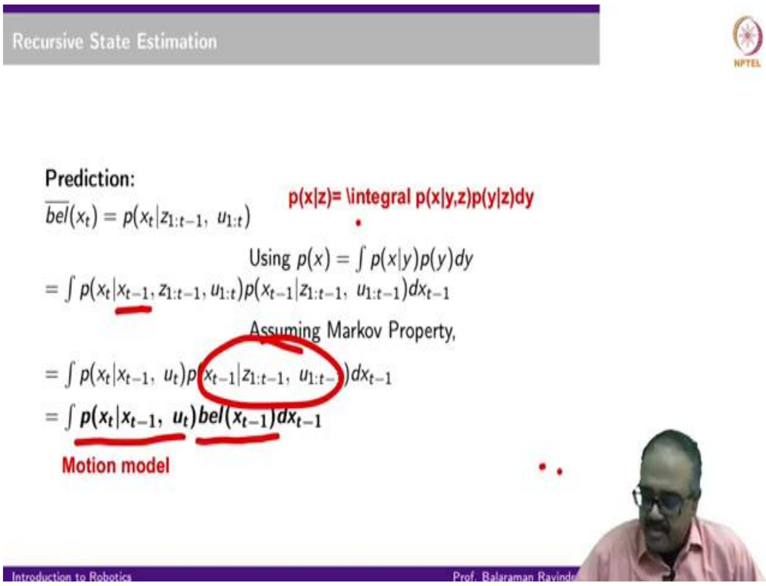

So that is a, so that is essentially our belief expression. So, I can, I can replace this with belief xt minus 1. So now my bel bar expression has now become something very simple. My bel bar xt equal to integral probability of xt given xt minus 1 comma ut. So, this is essentially my motion model, that is essentially my, so that is my motion model, so that is my motion model. And then I have my belief xt minus 1, dxt minus 1. So, my bel bar is essentially integral of the motion model times the previous belief, taken over all values that for the previous state. So that is basically it.

(Refer Slide Time: 22:10)

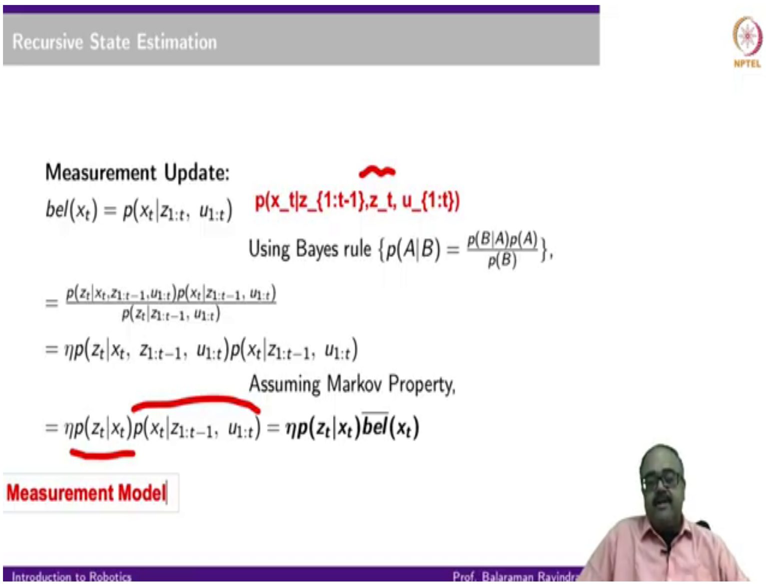

So now, once I have bel bar, now next thing would be to do the measurement update. So here in the measurement update, I go from bel bar, go from bel bar to bel. So, let us look at the bel definition again. So, bel xt is p of xt given z1 to t, u1 to t. So that is essentially what bel xt is.

Now I am going to use the Bayes rule, all of you are familiar with the Bayes rule, probability of A given B equal probability of B given A times probability of A divided by probability of B, I am pretty sure all of you are familiar with this rule.

Now, using that I am going to rewrite this. So, if you think of z1 to t, z1 to t is actually z1 to t minus 1 comma zt, so I am going to bring that out here. So now I am going to look at probability of zt, given xt comma z1 to t minus 1, u1 to t times probability of xt, given z1 to t minus 1, this is essentially in p of A part. So, I have, this is my p of B given A, this is my p of A, which is probably of xt, given all the previous measurements, and all the actions I have taken so far, and divided by probability of B part, which is probability of zt, given z1 to t minus 1 u1 to t.

So, this is essentially applying Bayes rule here assuming that I am looking at probability of xt, given z1 to t minus 1 comma zt comma u 1 to t. So, this is assuming that, that is my expression, as supposed to xt z1 to t. I am just splitting up z1 to t as z1 to t minus 1, zt and u1 to t, so I am assuming that my zt is my B and my xt is A and the rest of the variables are all conditioning variables. So therefore, I use the, I apply Bayes theorem and write it as probability of zt given xt, z1 to t minus 1, u1 to t times probability of xt given z1 to t minus 1 u1 to t divided by probability of B.

Now, if you think about it, I can again apply the Markov property because my xt has now gone to the conditioning part. I am assuming that I am, what is the, I am asking the question, what is the probability of zt, given xt. As soon as I ask that question, I can assume the Markov property. And I can simplify the first term in the expression as probability of zt, given xt, I do not have to worry about the history after that. So, this is what we said in the measurement model.

So, this guy is essentially my measurement model. And what about the rest of it, it is a probability of xt, given z1 to t minus 1 and u1 to t, and that is exactly my bel bar. So, this is my bel bar update, that is my bel bar function. So, I have it, I have my measurement model into my bel bar of xt. And the denominator actually does not depend on x itself and I can just replace it with some kind of a normalizing factor eta, which essentially, I could just sum the numerator for all possible values. And then normalize it so that I get a probability distribution.

I compute this P of bel of xt, you can think of some kind of temporary variable, for every possible value that xt can take, I compute the numerator, and then I divide all of these by the sum of the all the numerators I computed and that gives me the probability. So that is where the eta is, it is some kind of a normalizing factor. And I mean, technically eta is supposed to be probability of zt given z1 to t minus 1, u1 to t but it is a little hard to compute. And therefore, we just use this normalization trick.

So instead of actually computing the transition on the observations alone, without worrying about the states, I use known quantities, what are the known quantities I am using, I am using the measurement model. Let us go back. So, I am using the, so I am using the measurement model, and I am using bel bar, which we just computed on the previous slide using known quantities again. So, using both the measurement model and bel bar, we are able to compute what is belief without ever computing the denominator.

So, we just compute the numerators, and then normalize across all the possible values that xt could take. Now if xt could take continuous values, then normalization becomes a little tricky. And so, what we will see in the next few lectures are techniques for making this computation tractable, by making assumptions about the form of the transition function, the motion model, the form of the measurement model, and also the representation for the belief.

Right now, I have made no assumptions I am making, I am just telling you, these are very general probability distributions. As we go along, so we will start making specific assumptions about the transition model, we will start making specific assumptions about the measurement model and about the belief so that we come up with different kinds of tractable computations. Right now, eta is easy for you to compute, if you think that xt can take only a small number of finite values, small number of values.

But if xt can is like a continuous value output, so like orientation, you could just take any value or it could be velocity or even x y coordinates, like it could be anywhere in the workspace, then computing this eta, basically having to compete over all possible values for the belief is going to be a little tricky and we need to have some assumptions about how the belief looks like, so that I can do the computation tractably.