

PROFESSOR: Last time we talked about probability as an introduction on how to model uncertainty. State estimation is one of the ways that we can deal with uncertainty in our system. We can take a model that we don't completely understand and attempt to infer information about it based on the things that we can observe about that particular system.

In particular we're going to look at a set of observations and actions that either our system takes or that we observe about the system and that we take on that particular system. If we continue this process for multiple time steps then we can continue to attempt to learn things about the particular system. The process of completing that behavior over multiple time steps while making those inferences is what we refer to when we talk about state estimation.

First off, state estimation is a process that's completed as a consequence of wanting to understand a stochastic state machine. State estimation itself is not a stochastic state machine. State estimation attempts to take a stochastic state machine, make a model of that stochastic state machine, and then run state estimation on it iteratively to attempt to figure out, or recursively, an attempt to figure out what's going on inside that stochastic state machine.

When you build a stochastic state machine model there are three components that need to be specified. The first is the starting distribution over states. For instance, let's say that I believe that I am sick and I'm trying to figure out what it is that I am sick with. And I could be sick with three things, as far as I'm concerned. I could be sick with strep or I could be sick with some other more boring virus. Or I could be sick with mononucleosis.

The starting distribution refers to my starting belief as to the systems. And if I'm generically sick in the general sense, one of the assumptions that's frequently made with respect to starting distributions is that they're uniform, right? It could be equally any of these things.

The second thing you need to specify when you're talking about modeling a stochastic state machine is your observation distribution. Or what is the likelihood associated with making a particular observation given that you're in a current state? For instance, if I have mononucleosis how likely would it be that I observe a bunch of white ugly patches on the back of my throat? Or if I had strep? What is the likelihood associated with that? That kind of thing.

Typically this observation variable is factored into a couple different phenomena. In the sick example the best thing to talk about is symptoms. [INAUDIBLE] Am I lethargic? Do I have the white spots on the back of my throat? Do I have a fever? That sort of thing.

The last thing that you need specify when you're talking about modeling a stochastic state machine is your transition distribution. You assume that your state machine is going to change over time. Or it is likely that I will get more or less sick. And there are things that I can do to induce that kind of change. Or there are things that I can do that effectively model the passage of time.

Your actions for a stochastic state machine model can either be actions that the model takes and you were exclusively doing observations. But one of the particular observations that you do also qualifies as an action or something that indicates the passage of time. Or actions can be something that you do to a particular state.

In the sick example, things I could do to myself to try to make myself feel better or at least figure out better what is going on or what might cause my distribution to sway towards one particular state. I could take antibiotics. Or sleep in and drink a lot of orange juice. Or continue my day as normal.

Given a particular action, any particular state that you're starting from, your transition distribution tells you the likelihood associated with being in a new particular state. So as a consequence of making those actions, does the distribution of likelihood of a particular illness change?

At this point I'm going to walk through a step of state estimation. Each step of state

estimation is the same. In fact, if you complete multiple steps based on the information that you gained from the previous step, that's referred to as recursive state estimation. And I'll keep walking through the sick example.

So when you're doing state estimation you're trying to figure out something about a system that you cannot perfectly model. For instance, either your own immune system or your own susceptibility to a particular disease. And you have all the components you have for your stochastic state machine model.

As a consequence of the passage of time or as a consequence of making an observation and either observing an action taken by your stochastic state machine or performing an action upon your stochastic state machine you're going to make a new estimation of what you believe the current state of that unknown system. Or system that is not completely observable to you. You're going to make a new estimate of your belief of the state of that system.

In short you're going to solve for the probability distribution over $S_{(t+1)}$. There are two steps. The first step is referred to as the Bayesian reasoning step. And it involves performing Bayes evidence or Bayes rule upon the current state distribution given a particular observation.

So at this point I've made some sort of observation about myself. If I'm talking about the sick model, right? I spent all day coughing. Or I have a fever. Or my throat is sore. Or I feel extremely lethargic, right?

Given that observation I can take the $P(O \text{ given } S)$ from my observation distribution multiply it by my current understanding of the state distribution. And then divide out by $P(O)$.

The slowest way to complete this action is to build the joint distribution and then condition on a particular column. It's very proper. But you can save yourself some cycles by doing this.

Let's say I started off with the uniform distribution, right? It could be equally likely that I have strep or a normal virus or mono. As a consequence of making the

observation that I don't have white spots on the back of my throat. I could say, oh the likelihood of me being in that state-- the likelihood of me just having a normal virus is higher and the likelihood of me having either strep throat or mono is lower.

This step takes $P(S)$ and multiplies it by $P(O \text{ given } S)$. Once I have these values I have to scope back out to the universe or I have to normalize these values such that they sum to 1.

That's where I get my $P(S_t \text{ given } O)$. At this point I've accounted for the observation that I've made, but I haven't accounted for the action on the system. That's the next step.

We're going to take our results of Bayesian reasoning which are sometimes referred to as $B \text{ prime } S_t$. And take the action and find the distribution overstates as a consequence of a single time step or a single iteration of state estimation.

The second step is referred to as a transition update. We've got our updated belief. We're going to take our transition distribution or our specification for what happens given that we're in a current state and an action has been taken. At that point we'll have a probability distribution over the new states. And here are my values from the first step.

As an example let's say that I sleep in and drink a lot of orange juice. As a consequence of sleeping in and drinking a lot of orange juice there's some amount of likelihood that I will either continue to be sick with strep or it's possible that I actually have just a normal virus. If I have a normal virus and I sleep in and drink a lot of orange juice, this causality sounds backwards but it's as a consequence of not being able to make perfect observations on the system.

If I have an amount of belief that says that I think I have strep and I sleep in and drink a lot of orange juice, then it's equally likely that I will have either strep or a normal virus after completing that step, right? It doesn't differentiate between the two.

If I'm sick with a virus and I sleep in and drink a lot of orange juice, then the state

that I'm going to encourage myself to be in is I have a virus. If I have mono and I sleep in and drink a lot of orange juice then there's some likelihood on the next day that I will still be in a state that looks like I have mono. But there's also some likelihood associated with it that I will be in some state that looks like I have strep. That's what happens when you run the transition update.

When you run the transition update you end up accumulating all the probabilities associated with being in a particular new state. As a consequence of being in a particular previous state and entering that new state based on the transition distribution. Once you accumulate all these values you end up with your new distribution over a new state.

This represents one step of state estimation. If I wanted to run multiple, I would take the value that I got here for S_{t+1} , replace it in the value for S_t . And run the same process of Bayesian reasoning and transition update. This concludes my review of state estimation. Next time we'll talk about search.