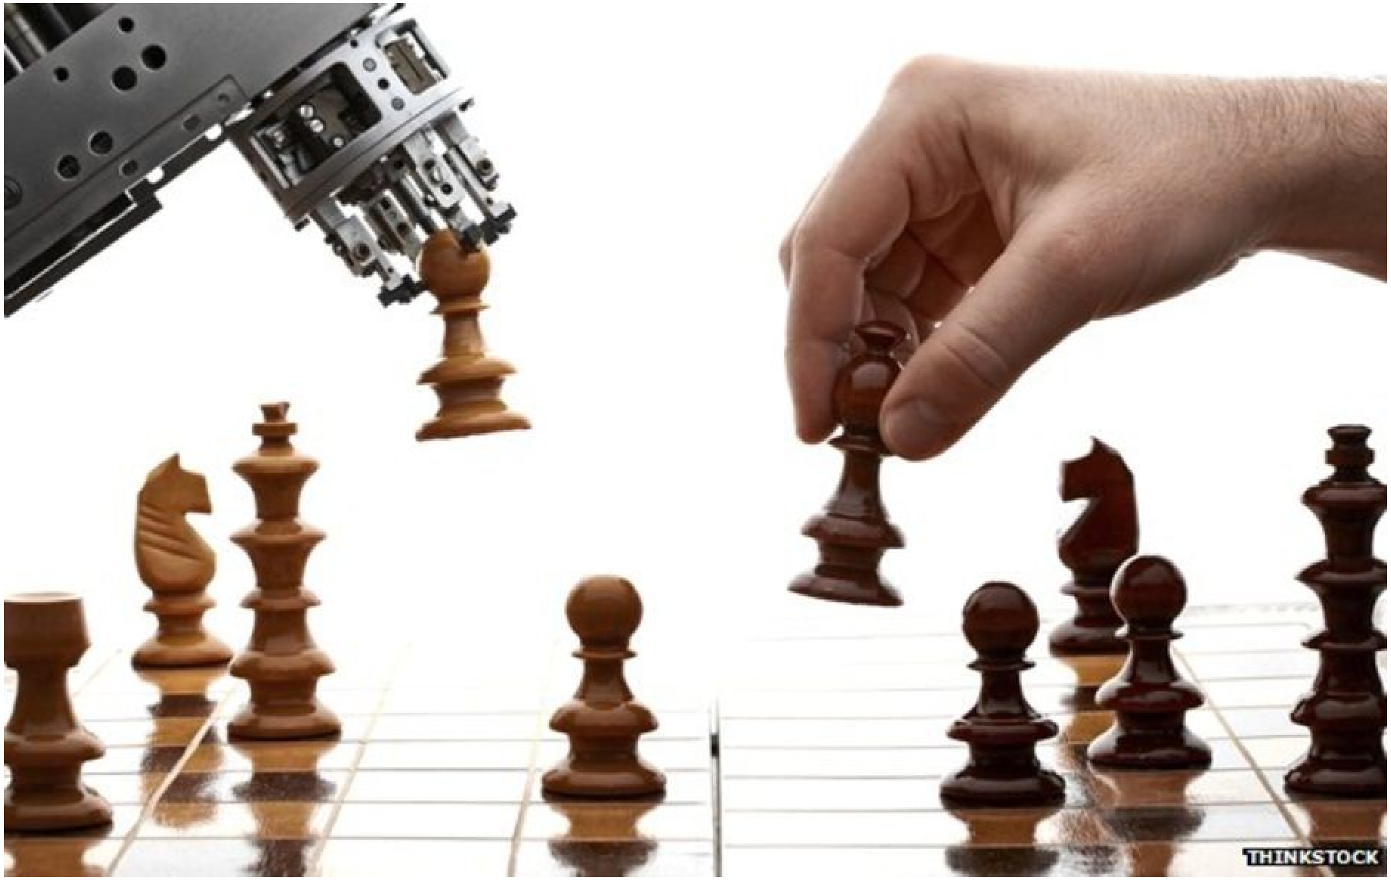

Introduction to Artificial Intelligence

COSC 4550 / COSC 5550

Professor Cheney

10/2/17

instead, let's assume that all the information necessary to know what state X_t will be is contained in state X_{t-1}

(i.e. conditionally independent of $X_{0:t-2}$)

$$P(X_t \mid X_{t-1}, X_{t-2}, X_{t-3}, \ldots, X_1, X_0) = P(X_t \mid X_{t-1})$$

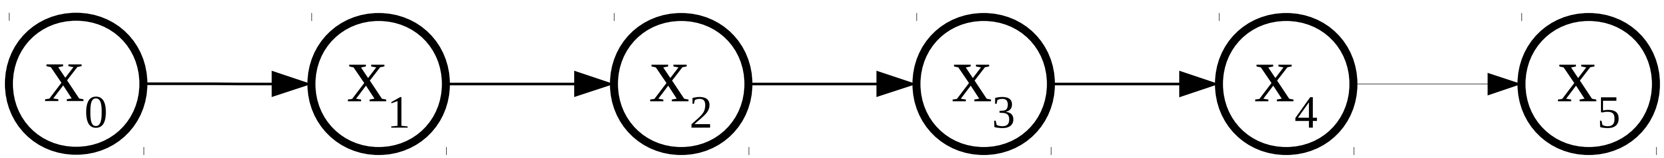

“Markov assumption”

(resulting process is called a “Markov process” or “Markov chain”)

“sensor Markov assumption”

we'll make the Markov assumption about both observable evidence variables (E) and non-observable hidden state variables (X)

$$P(E_t \mid X_{0:t-1}, E_{0:t-1}) = P(E_t \mid X_{t-1})$$

- $P(E_t$: estimates of your current sensor values
- $X_{0:t-1}, E_{0:t-1}$: informed by *all* your past sensors and past (hidden) world states
- $=$: are no better than
- $P(E_t$: an estimate of your current sensor values
- X_{t-1}: informed only by the (hidden) world state at the *last* timestep

e.g. you're a security guard that works in the basement
you might not be able to observe the weather directly,
but you can make a prediction about it each day based on
whether or not your boss is holding an umbrella when arrives

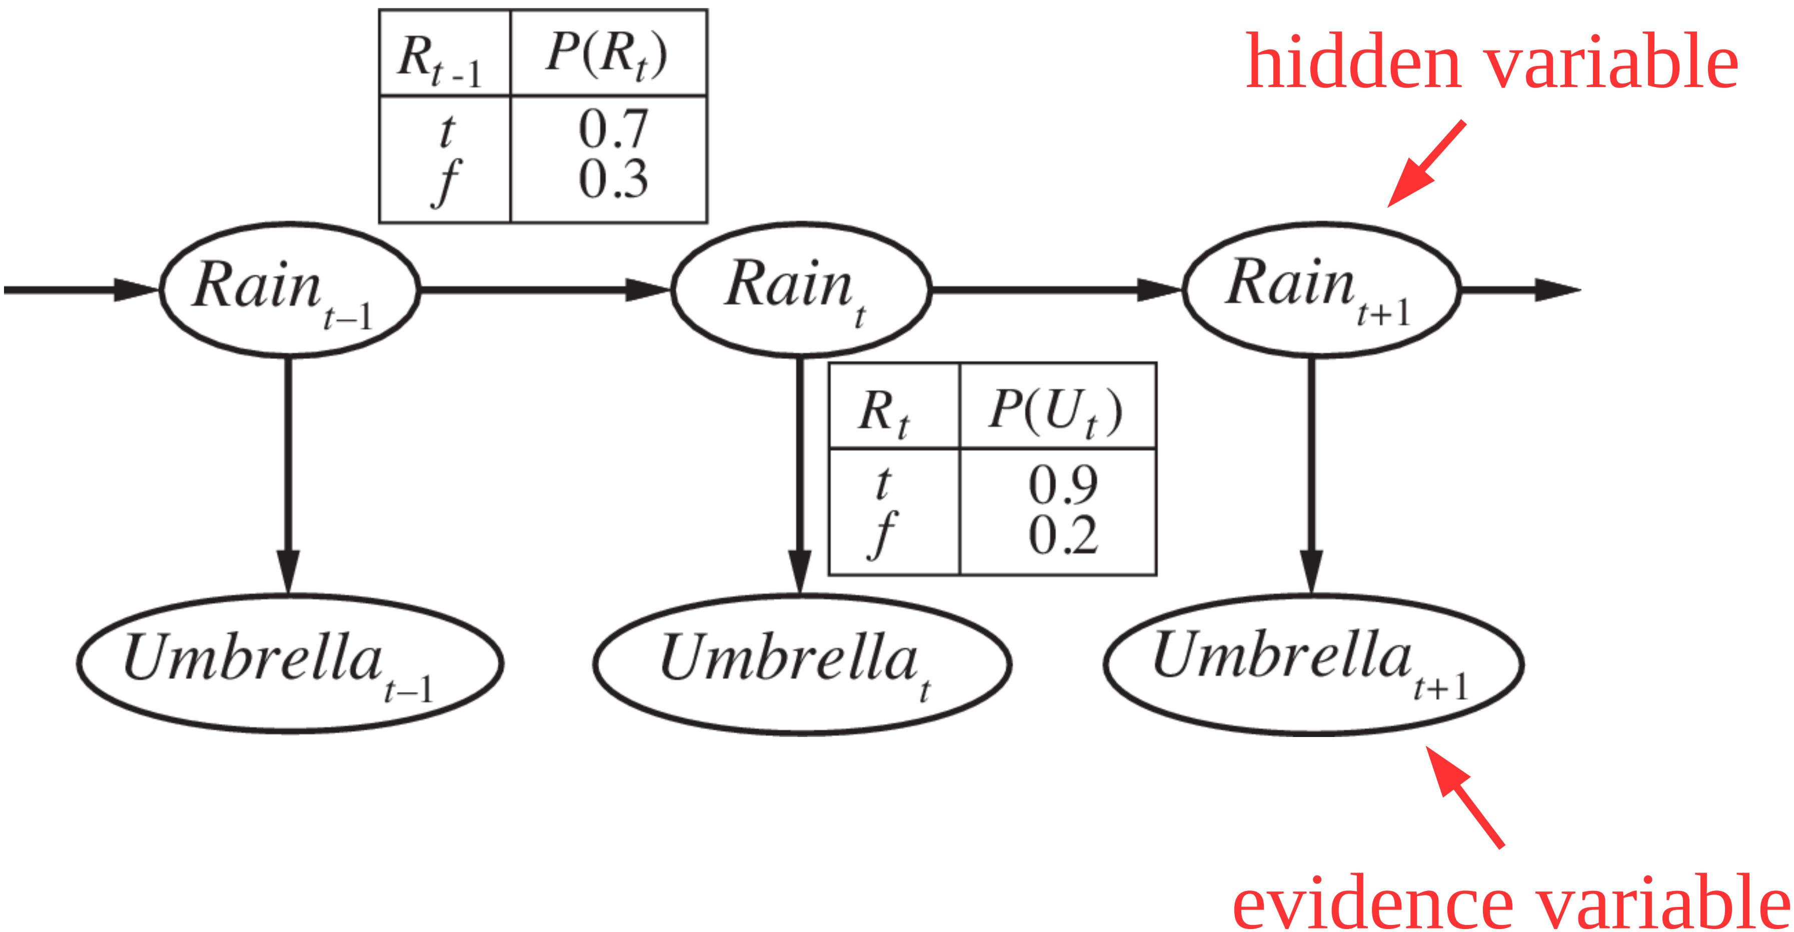

e.g. you're a security guard that works in the basement
you might not be able to observe the weather directly,
but you can make a prediction about it each day based on
whether or not your boss is holding an umbrella when arrives

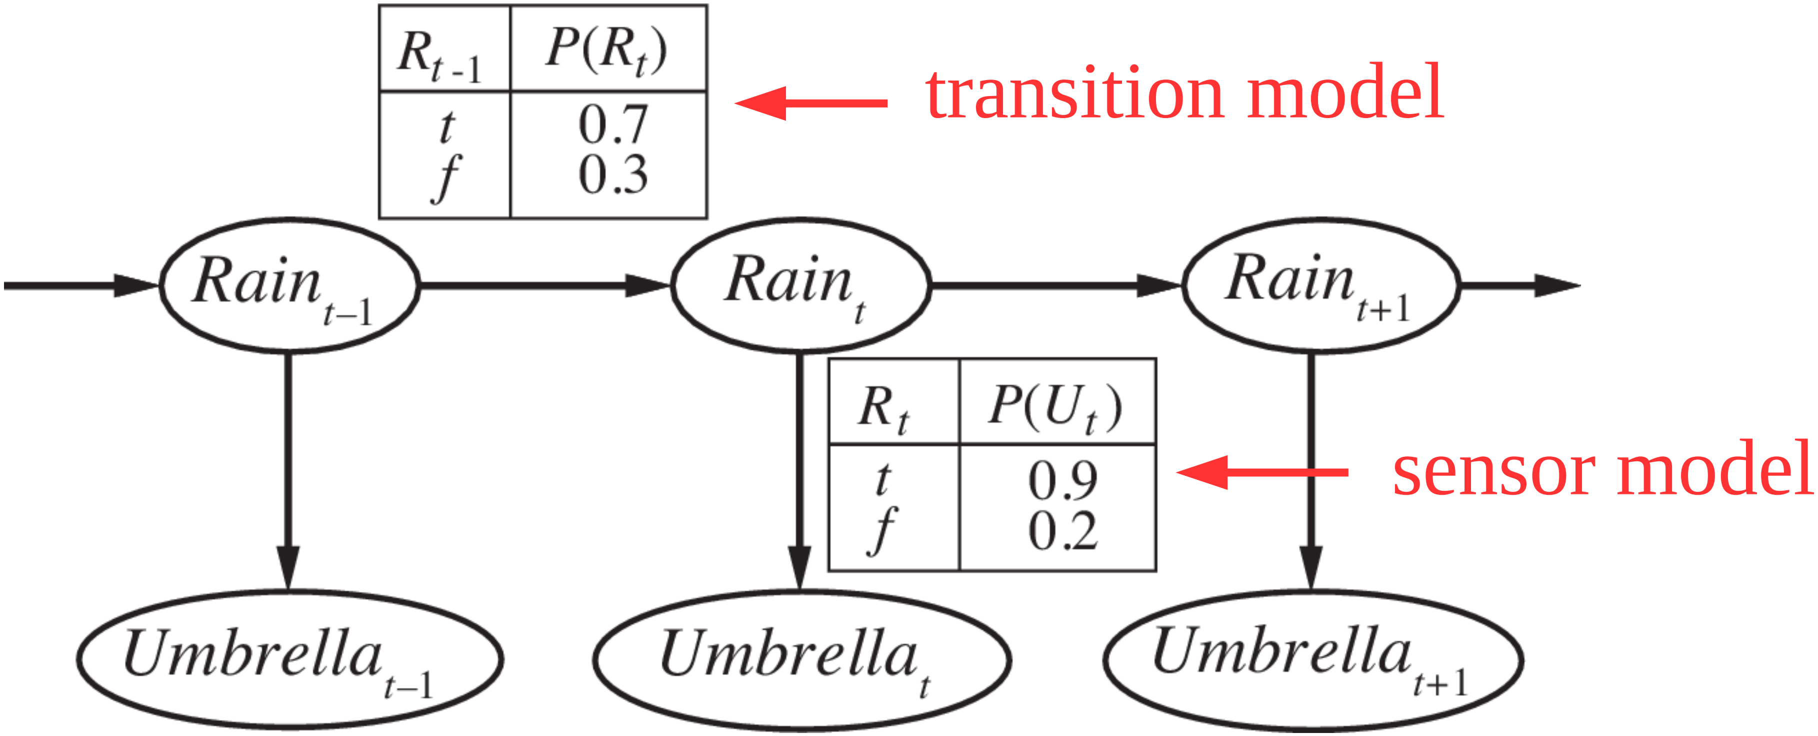

note: even though the world is changing every timestep,
the *way* in which it changes (the transition model)
is assumed to hold constant throughout learning/use

for each timestep

$$\mathbf{P}(\mathbf{X}_{0:t}, \mathbf{E}_{1:t}) = \mathbf{P}(\mathbf{X}_0) \prod_{i=1}^{t} \mathbf{P}(\mathbf{X}_i \mid \mathbf{X}_{i-1})\, \mathbf{P}(\mathbf{E}_i \mid \mathbf{X}_i)$$

prior — transition model — sensor model

$P(X_0, E_0) = P(E_0 | X_0) * P(X_0)$

$P(X_1, E_1) = P(E_1 | X_1) * P(X_1 | X_0) * P(X_0)$

$P(X_2, E_2) = P(E_2 | X_2) * P(X_2 | X_1) * P(X_1 | X_0) * P(X_0)$

$P(X_3, E_3) = P(E_3 | X_3) * P(X_3 | X_2) * P(X_2 | X_1) * P(X_1 | X_0) * P(X_0)$

...

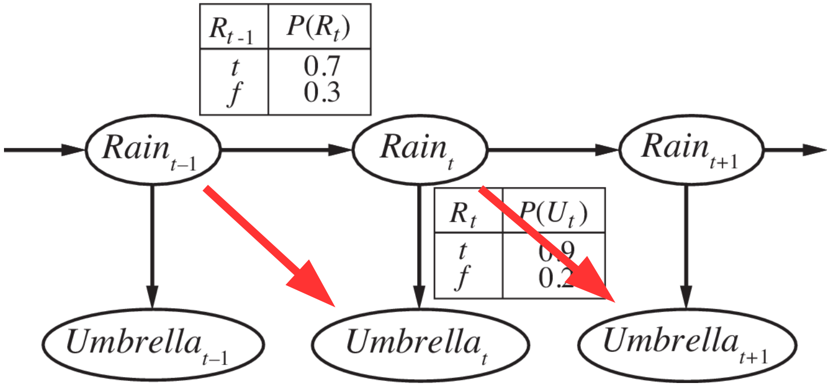

is the Markov assumption reasonable here?

how could we improve it?

(1) look more than one step behind us

$$P(Umbrella_t, Rain_t \mid Umbrella_{0:t-1}, Rain_{0:t-1}) = P(Umbrella_t, Rain_t \mid Rain_{t-1}, Rain_{t-2})$$

(second order Markov process)
(can also do third order, fourth order, ….)

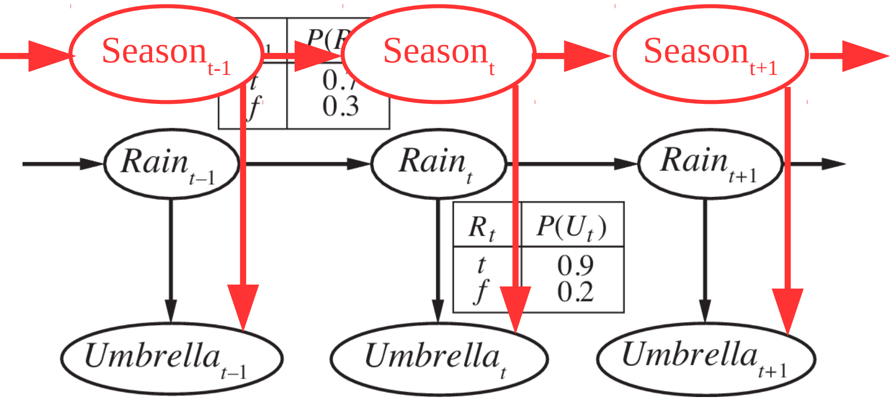

is the Markov assumption reasonable here?

how could we improve it?

(2) add more world state variables

$$P(\text{Umbrella}_t, \text{Rain}_t \mid \text{Umbrella}_{0:t-1}, \text{Rain}_{0:t-1}) = P(\text{Umbrella}_t, \text{Rain}_t \mid \text{Rain}_{t-1}, \text{Season}_{t-1})$$

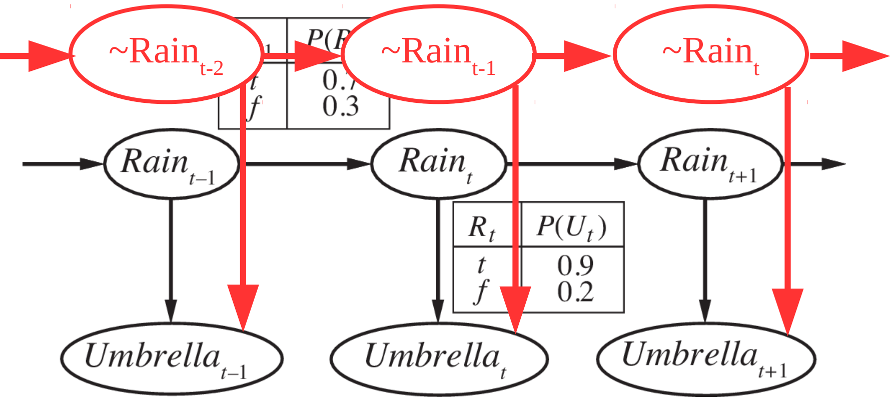

is the Markov assumption reasonable here?

how could we improve it?

(2) add more world state variables

$$P(Umbrella_t, Rain_t \mid Umbrella_{0:t-1}, Rain_{0:t-1}) = P(Umbrella_t, Rain_t \mid Rain_{t-1}, Season_{t-1})$$

note: option (1) is a subset of (2)
(i.e. your additional variable could contain information about past states, such that any process can be defined as a first-order Markov decision process, with enough hidden states)

we now have our temporal inference model

what fun/useful stuff can we do with it?

filtering
prediction
smoothing
most-likely explanation

filtering

“belief state estimation”

“filter” down all the observed evidence to
determine what the hidden state variables are
(i.e. filter down to only the features that drive our transitions)
and tell me the current hidden state

$$P(\ X_t \mid e_{1:t}\ )$$

e.g. what is $P(R_2 \mid u_{1,2})$?

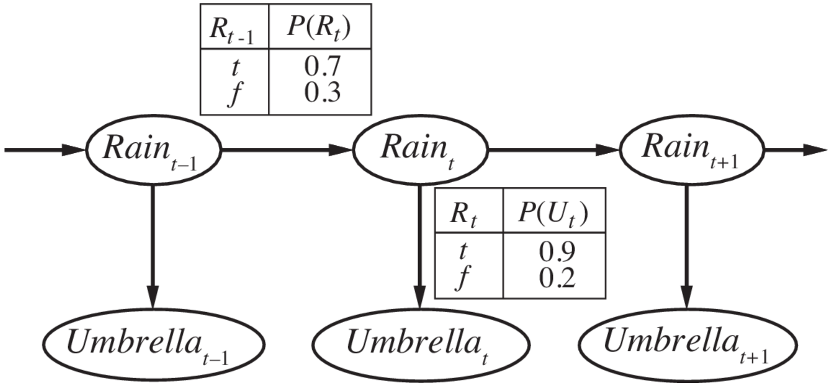

$P(R_0)$ = [0.5, 0.5] (let's assume no prior knowledge/bias)

$P(R_1) = \sum_{r0} P(R_1 \mid r_0) * P(r_0)$ ← chain rule

$= [0.7, 0.3] * 0.5 + [0.3, 0.7] * 0.5$

$= [0.5, 0.5]$

(we didn't observe anything,
so we didn't learn anything)

$P(R_1) = [0.5, 0.5]$

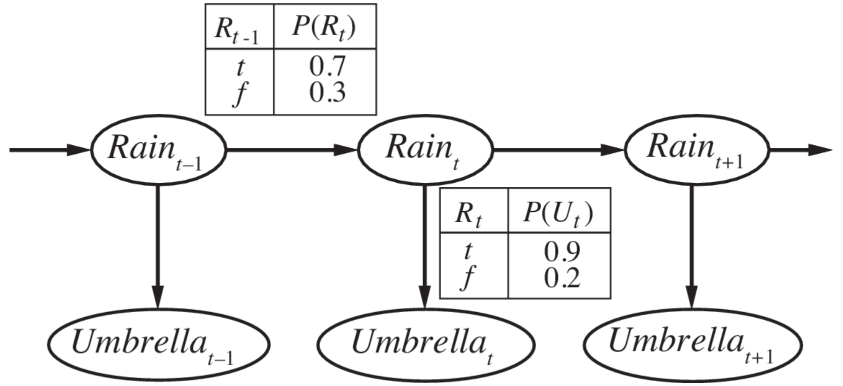

on day 1, we observe that u_1= true
(boss brought his umbrella)

$P(R_1 \mid u_1) = \alpha\, P(u_1 \mid R_1) * P(R_1)$ ← Bayes' rule!
$= \alpha\, [0.9, 0.2] * [0.5, 0.5]$
$= \alpha\, [0.45, 0.1] \approx [0.818, 0.812]$

$P(R_2 \mid u_1) = \sum_{r1} P(R_2 \mid r_1) * P(r_1 \mid u_1)$ ← chain rule
$= [0.7, 0.3] * 0.818 + [0.3, 0.7] * 0.182$
$= [0.627, 0.373]$

$P(R_2 \mid u_1) = [0.627, 0.373]$

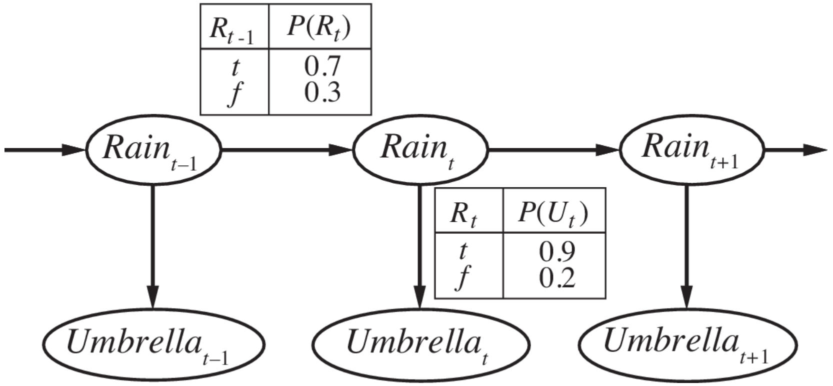

on day 2, we observe that u_2= true
(boss brought his umbrella again)

$$P(R_2 \mid u_1, u_2) = \alpha\, P(u_2 \mid R_2) * P(R_2 \mid u_1)$$
$$= \alpha\, [0.9, 0.2] * [0.627, 0.373]$$
$$= \alpha\, [0.565, 0.075] \approx [0.883, 0.117]$$

Bayes' rule! (and Markov assumption)

we've found the current hidden (world) state from our history of observed states! ("filtering")

prediction

what will the world look like next?

$$P(X_{t+k} \mid e_{1:t})$$

for $k > 0$

how do you do it?

we already did it!!!

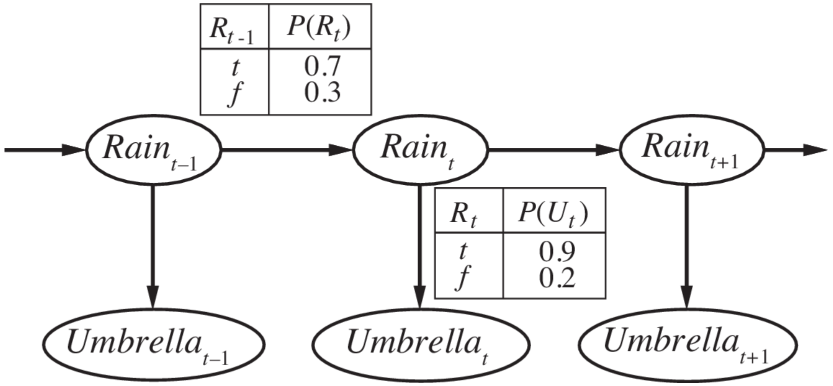

$$P(R_1) = [0.5, 0.5]$$

$$P(R_1 \mid u_1) = \alpha\, P(u_1 \mid R_1) * P(R_1)$$ ← Bayes' rule!

$$= \alpha\, [0.9, 0.2] * [0.5, 0.5]$$

$$= \alpha\, [0.45, 0.1] \approx [0.818, 0.182]$$

$$P(R_2 \mid u_1) = \sum_{r1} P(R_2 \mid r_1) * P(r_1 \mid u_1)$$ ← chain rule

$$= [0.7, 0.3] * 0.818 + [0.3, 0.7] * 0.182$$

$$= [0.627, 0.373]$$

↑ prediction!

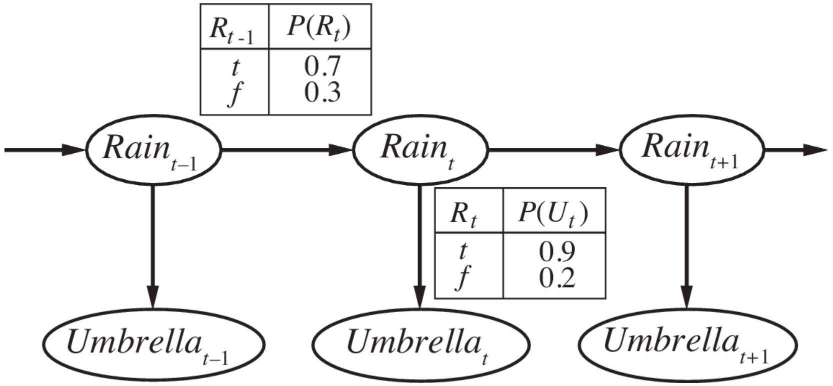

$$P(R_2 \mid u_1) = \sum_{r1} P(R_2 \mid r_1) * P(r_1 \mid u_1)$$
$$= [0.7, 0.3] * 0.818 + [0.3, 0.7] * 0.182$$
$$= [0.627, 0.373]$$

prediction!

note how our forward prediction made us less certain than we were are the previous time step (i.e. closer to [0.5, 0.5], e.g. [0.818, 0.182] → [0.627, 0.373])

if we predict far enough into the future, our observations are no longer relevant (and we converge to [0.5, 0.5])

smoothing

given what we know now (many observations later),
can we go back and update our guess of a past hidden state?

$$P(X_k \mid e_{1:t})$$

for $0 < k < t$

for this we need both forward inference,
and backwards inference

let's go back and smooth our guess of rain on day 1, given that we saw an umbrella on day 2

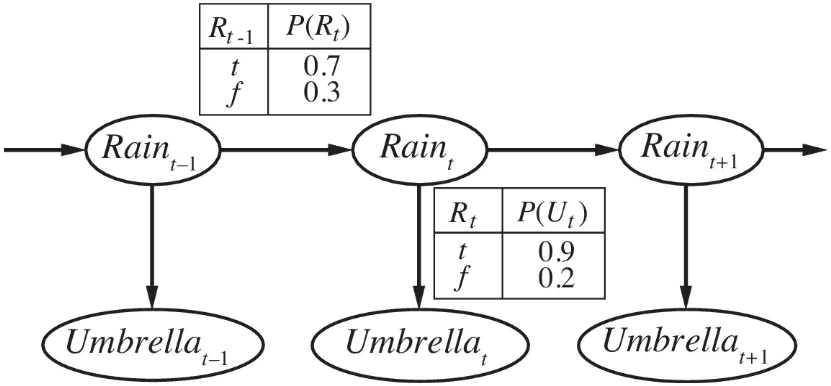

filtering = forward   backward

$$P(R_1 \mid u_{1,2}) = \alpha\, P(R_1 \mid u_1) * P(u_2 \mid R_1)$$
$$= \alpha\, [0.818, 0.182] * \sum_{r2} P(u_2 \mid r_2) * P(r_2 \mid R_1)$$
$$= \alpha\, [0.818, 0.182] * (0.9*[0.7, 0.3]+0.2*[0.3, 0.7])$$
$$= \alpha\, [0.564, 0.075] \approx [0.883, 0.117]$$

(note: prediction improved from when we only had forward pass)
$[0.818, 0.182] \rightarrow [0.883, 0.117]$

"forward-backward algorithm"

most-likely sequence

given a sequence of observation in time,
what's the most likely set of hidden states to cause them?

$$\max_{x1,\, x2,\, \ldots,\, xt} P(x_1, x_2, \ldots, x_t, X_{t+1} \mid e_{1:t+1})$$

e.g. given a set of sounds, what the most likely sequence of words that the user was trying to say?

e.g. given a set of umbrella observations,
what's the most likely weather to have caused it?

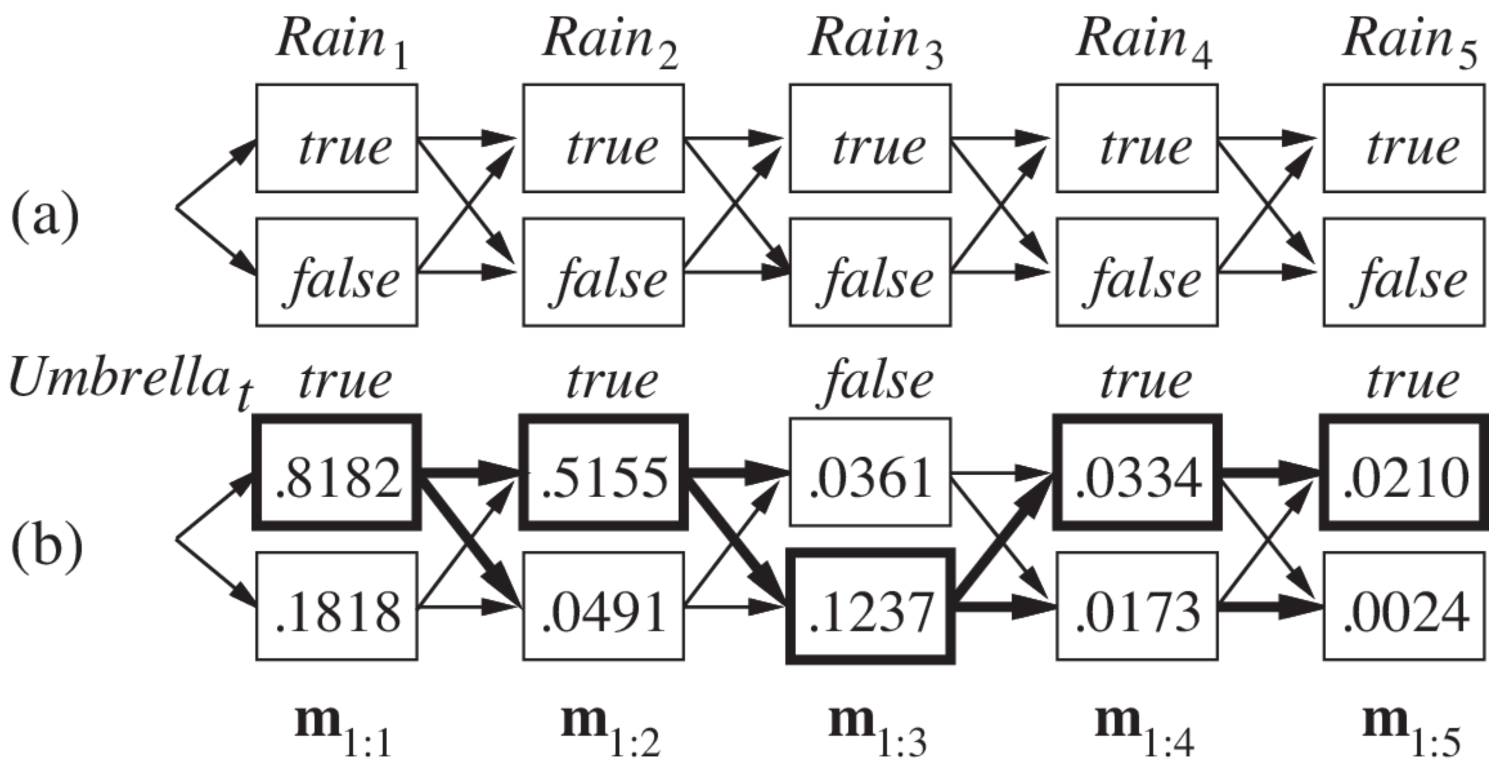

similar to filtering (i.e. finding next hidden state) over time
"Viterbi algorithm" (in your book)