CSC421 Intro to Artificial Intelligence

UNIT 25:
Probabilistic Reasoning over Time
(+ approximate inf. leftovers)

A little sidenote on Financial Engineering

- Financial instruments
  - Ways to manage risk
- Example: livestock producer purchase futures (type of instrument) to cover feed costs, so that they can plan on a fixed cost of feed. Typically the price is higher than the current market price
- If in future price drops bank is profitable if price goes up banks might loose money
- Question: how much to charge extra for the risk protection ? How many stocks of feed companies to buy/sell ?

Inference by stochastic simulation

- Basic idea
  - Draw N samples from a sampling distribution S
  - Compute an approximate posterior P'
  - Show this converges to true prob. P

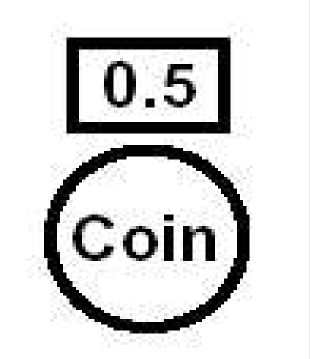

- Outline
  - Sampling from an empty network
  - Rejection sampling: reject samples disagreeing with evidence
  - Likelihood weighting: use evidence to weight samples
  - Markov Chain Monte Carlo (MCMC): sample from a stochastic process whose stationary distribution is the true posterior

Likelihood Weighting

- Idea: fix evidence variables, sample only nonevidence variables and weight each sample with the likelihood it affords the evidence
- Produces consistent estimates (details in book)

Example

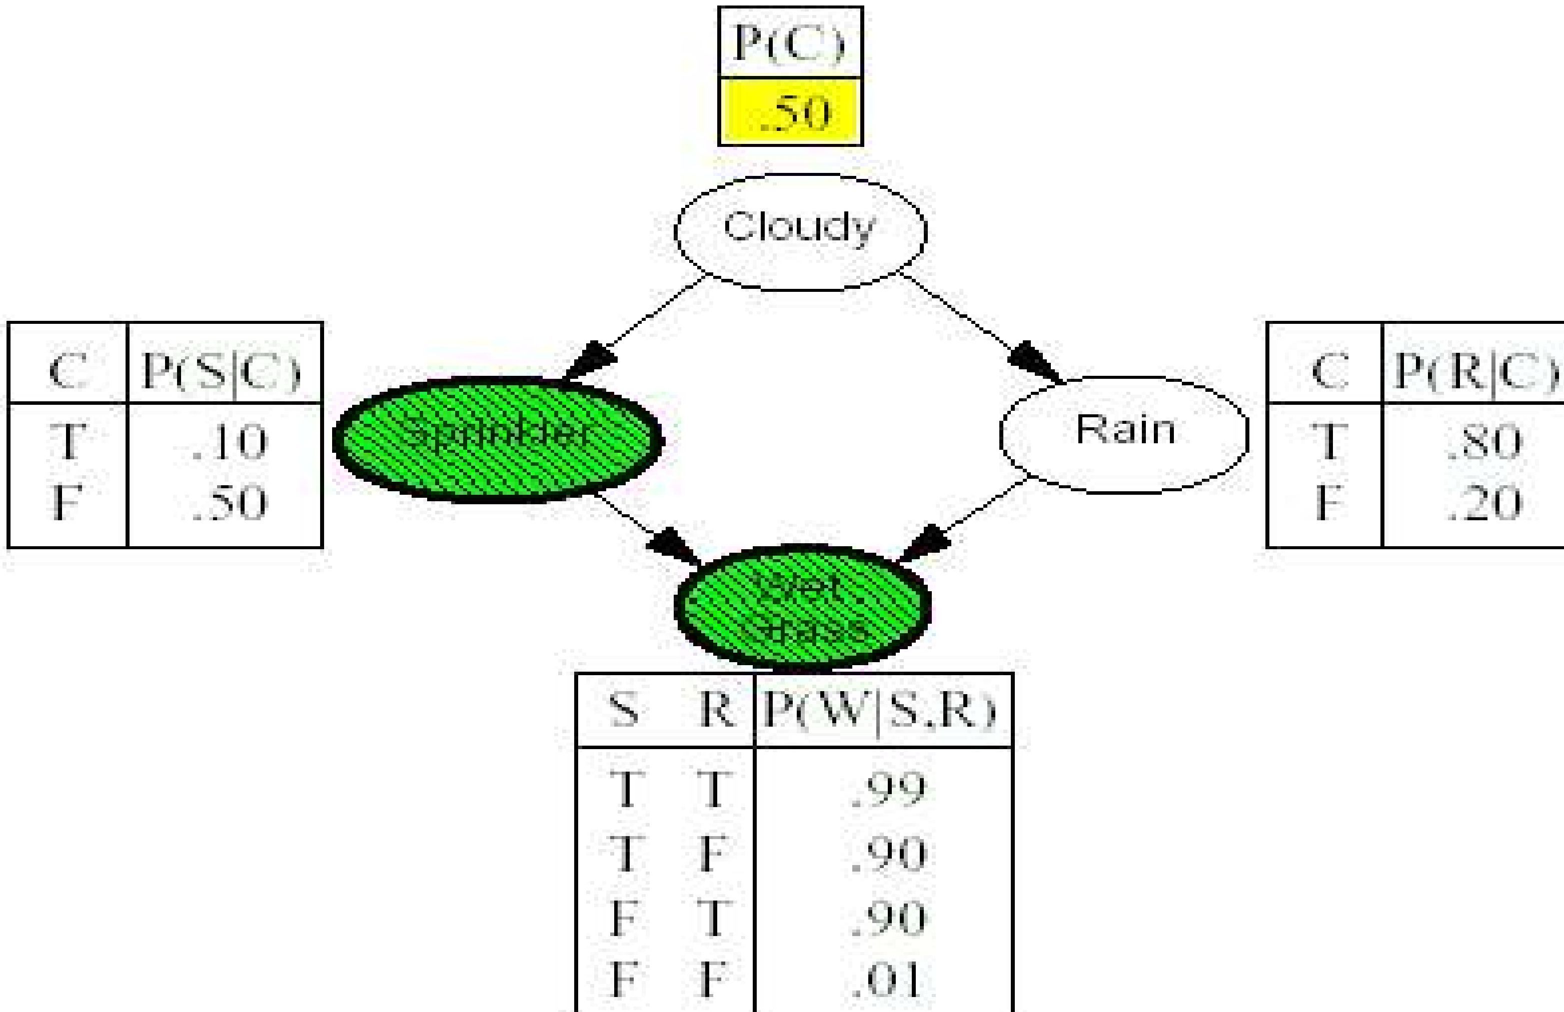

P(C)
.50

C	P(S|C)
T	.10
F	.50

C	P(R|C)
T	.80
F	.20

S	R	P(W|S,R)
T	T	.99
T	F	.90
F	T	.90
F	F	.01

$w = 1.0$

Example

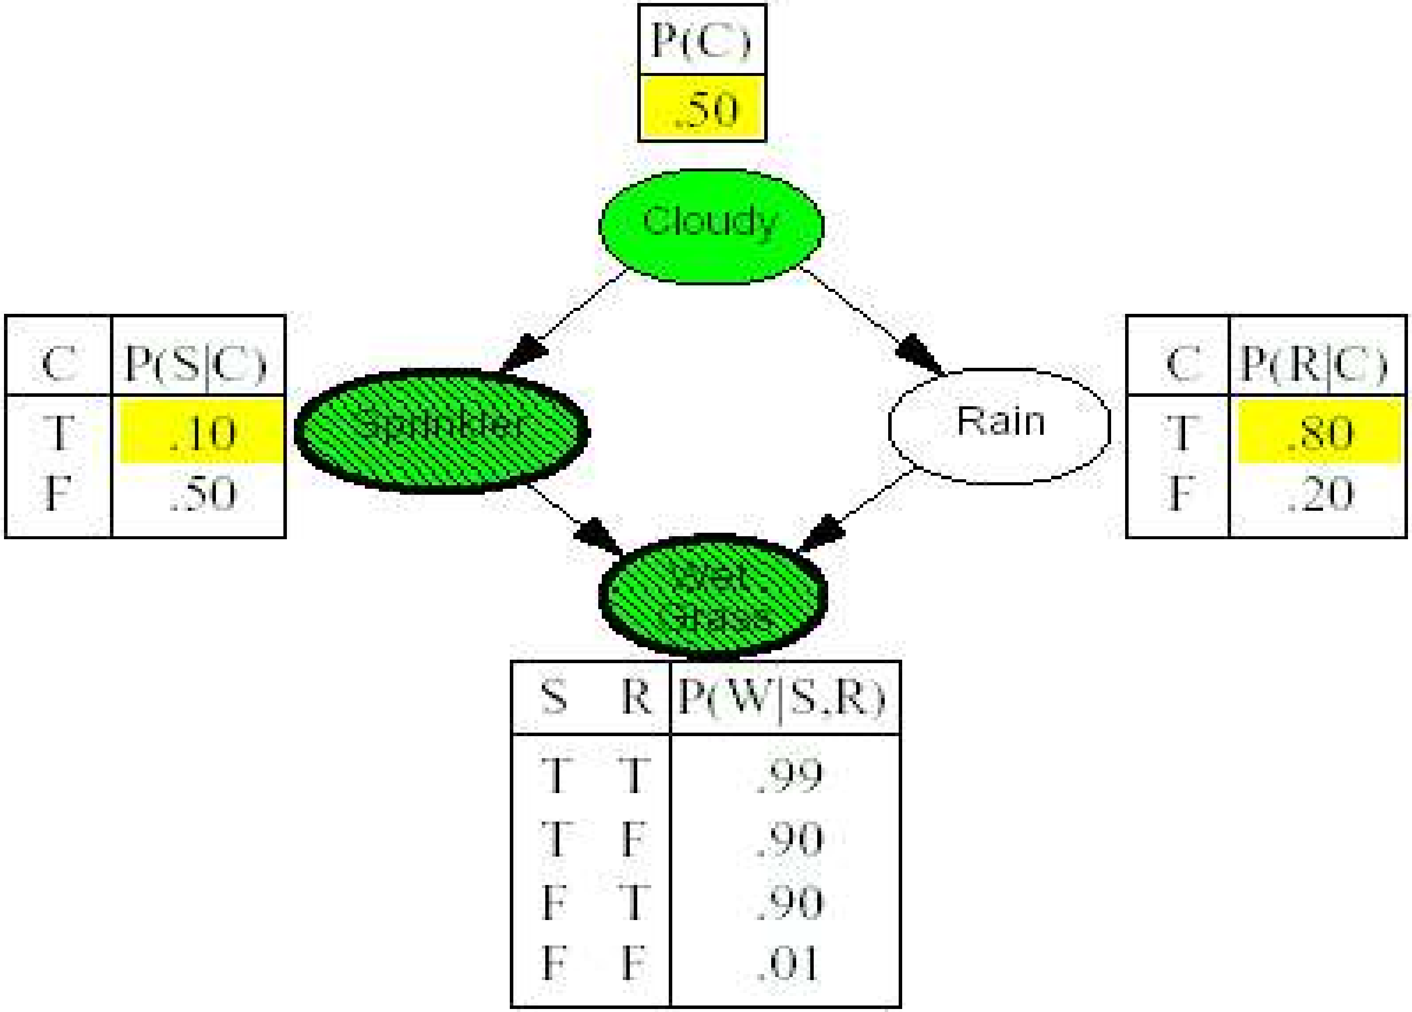

P(C)
.50

C	P(S|C)
T	.10
F	.50

C	P(R|C)
T	.80
F	.20

S	R	P(W|S,R)
T	T	.99
T	F	.90
F	T	.90
F	F	.01

$w = 1.0 \times 0.1$

Example

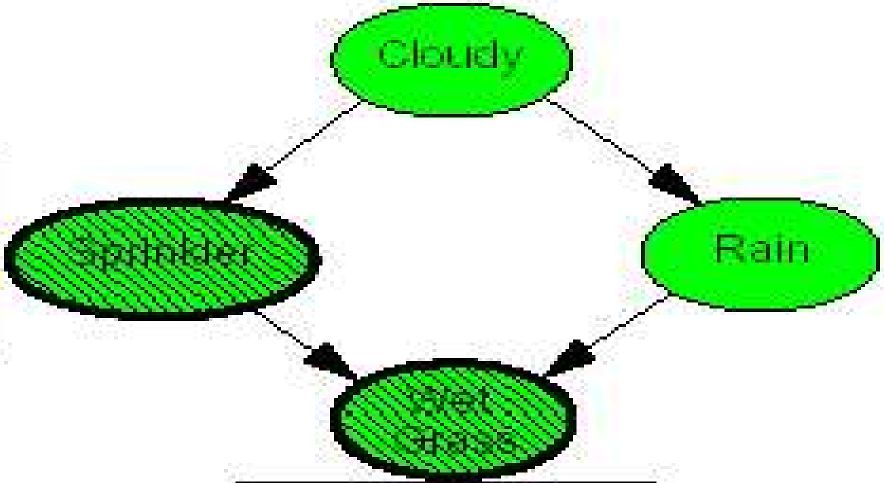

P(C)
.50

C	P(S|C)
T	.10
F	.50

C	P(R|C)
T	.80
F	.20

S	R	P(W|S,R)
T	T	.99
T	F	.90
F	T	.90
F	F	.01

$$w = 1.0 \times 0.1 \times 0.99 = 0.099$$

Approximate inference using MCMC

- State of network = current assignment of all network variables
- Generate next state by sampling one variable given Markov Blanket
- Sample each variable in turn keeping evidence fixed
- Can also choose a variable to sample at random each time

Markov Chain

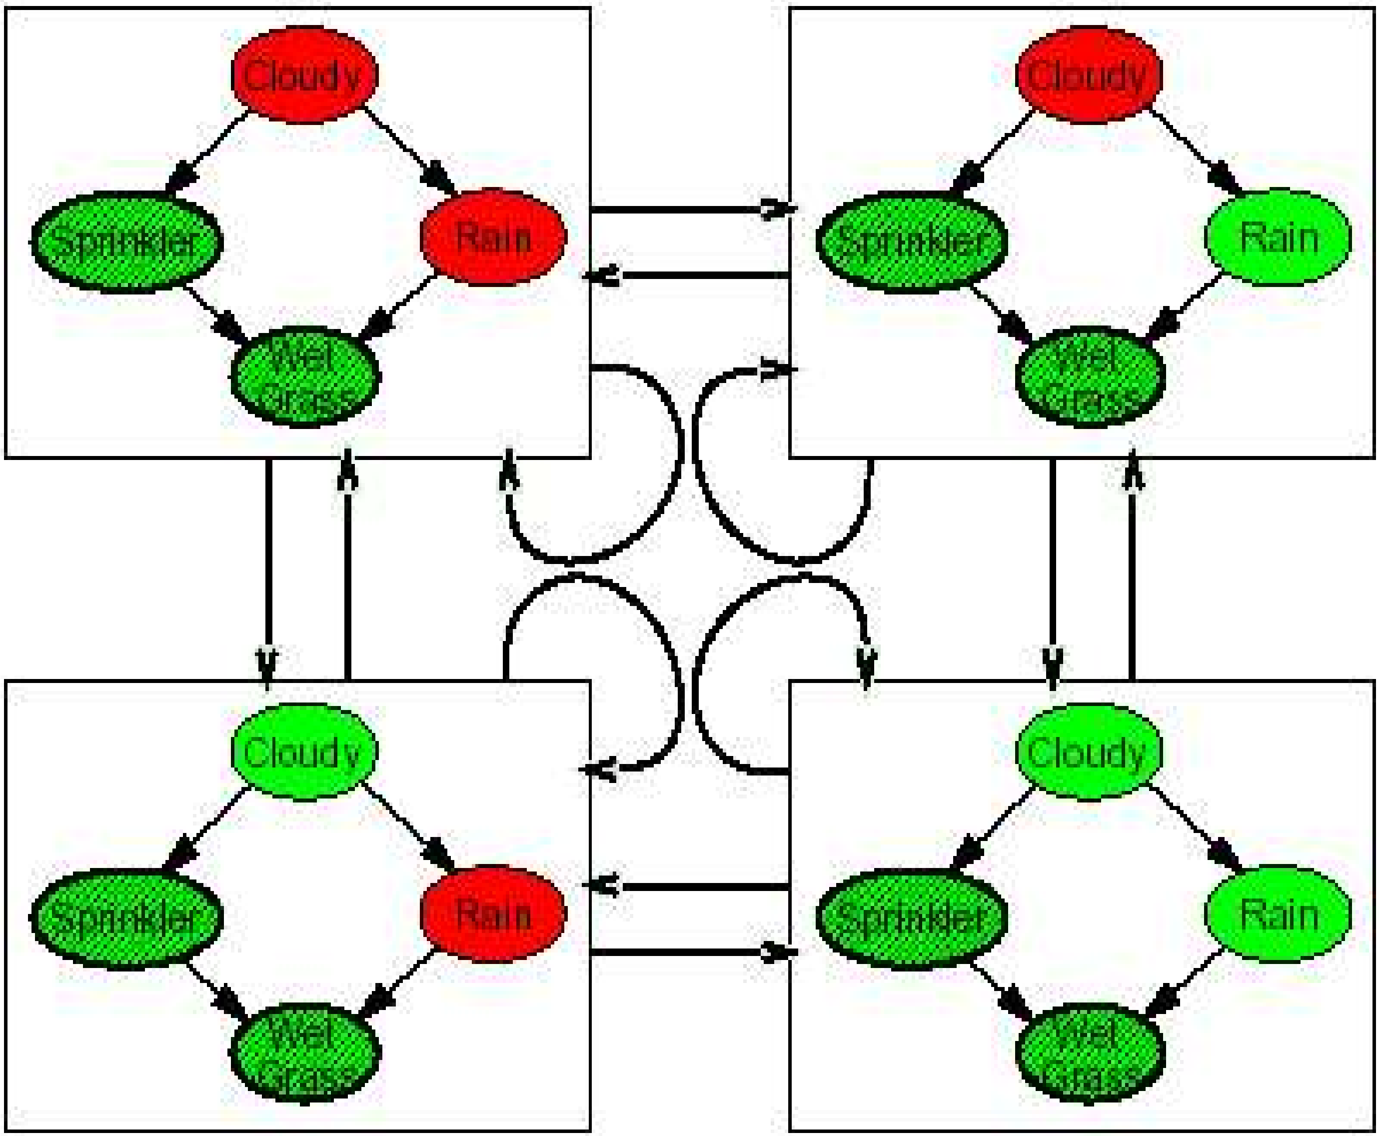

Wander around a bit average what you see

Transition probabilities

Every time a variable is sampled new state

Markov Blanket

- Markov Blanket = the parents, children and children's parents
- Example:
  - P(Rain | Sprinkler = true, WetGrass = true)
  - Cloudy , Rainy initialized randomly [t,f]
  - [t, t, f, t]
- Cloudy is sampled from P(Cloudy | s, r). Suppose Cloudy = false. Then the new current state is [f, t, f,t].
- Repeat wandering around
- Each state is a sample

Why this works ?

- Theorem:
  - Chain approaches stationary distribution: long run fraction of time spent in each state is exactly proportional to its posterior probability

Markov Blanket Sampling

- Markov Blanket of Cloudy is Sprinkler and Rain
- Markov Blanket of Rain is Cloudy, Sprinkler,and Wet Grass
- Prob. given Markov blanket = $P(x \mid MB(X)) = P(x \mid Parents(X)) * \prod_{z \in children(X)} P(z \mid parents(Z))$

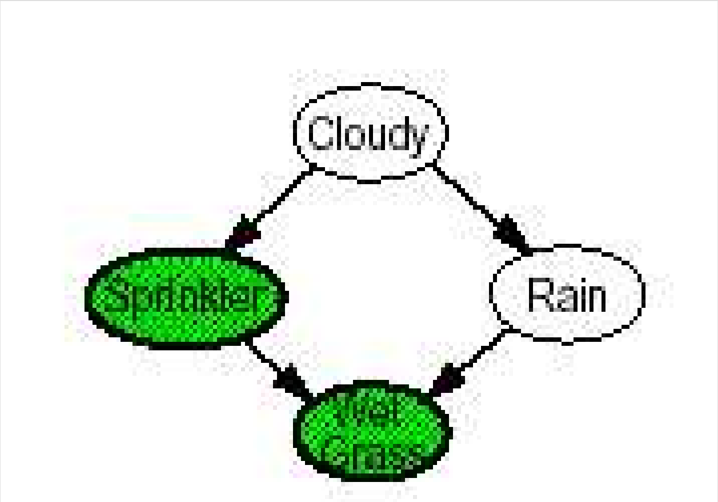

Summary

- Exact inference by variable elimination
  - Polytime on polytrees, NP-hard on general graphs
  - Space = time, very sensitive to topology
- Approximate inference by LW, MCMC
  - LW, MCMC generally insensitive to topology
  - Convergence can be slow with probabilities close to 0 or 1
  - Can handle arbitrary combinations of discrete and continuous random variables
  - USED HEAVILY IN PRACTICE

Time and Uncertainty

- The world changes; we need to track & predict
- Diabetes management vs vehicle diagnosis
- Basic idea: copy state + evidence for each time step
- X_t set of unobservable state variables at t
  - e.g. BloodSugar, StomachContents
- Y_t set of observable evidence variables at t
  - e.g MeasuredBloodSugar, PulseRate, FoodEaten
- Assumes discrete time
- Notation: $X_{a:b} = X_a, X_{a+1}, \ldots, X_{b-1}, X_b$

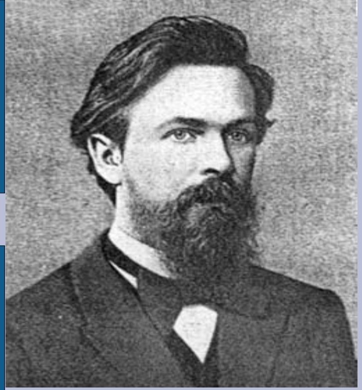

Markov Processes (or chains)

- Markov assumption
  - X_t depends on bounded subset of $X_{0:t-1}$
  - First-order: $P(X_t | X_{0:t-1}) = P(X_t | X_{t-1})$
  - Second-order: $P(X_t | X_{0:t-1}) = P(X_t | X_{t-2}, X_{t-1})$
  - ...
- Sensor Markov Assumption
  - $P(E_t | X_{0:t}, E_{0:t-1}) = P(E_t | X_t)$
- Stationary process
  - Transition model $P(X_t | X_{t-t})$ fixed for all t
  - Sensory model $P(E_t | X_t)$ fixed for all t

Example

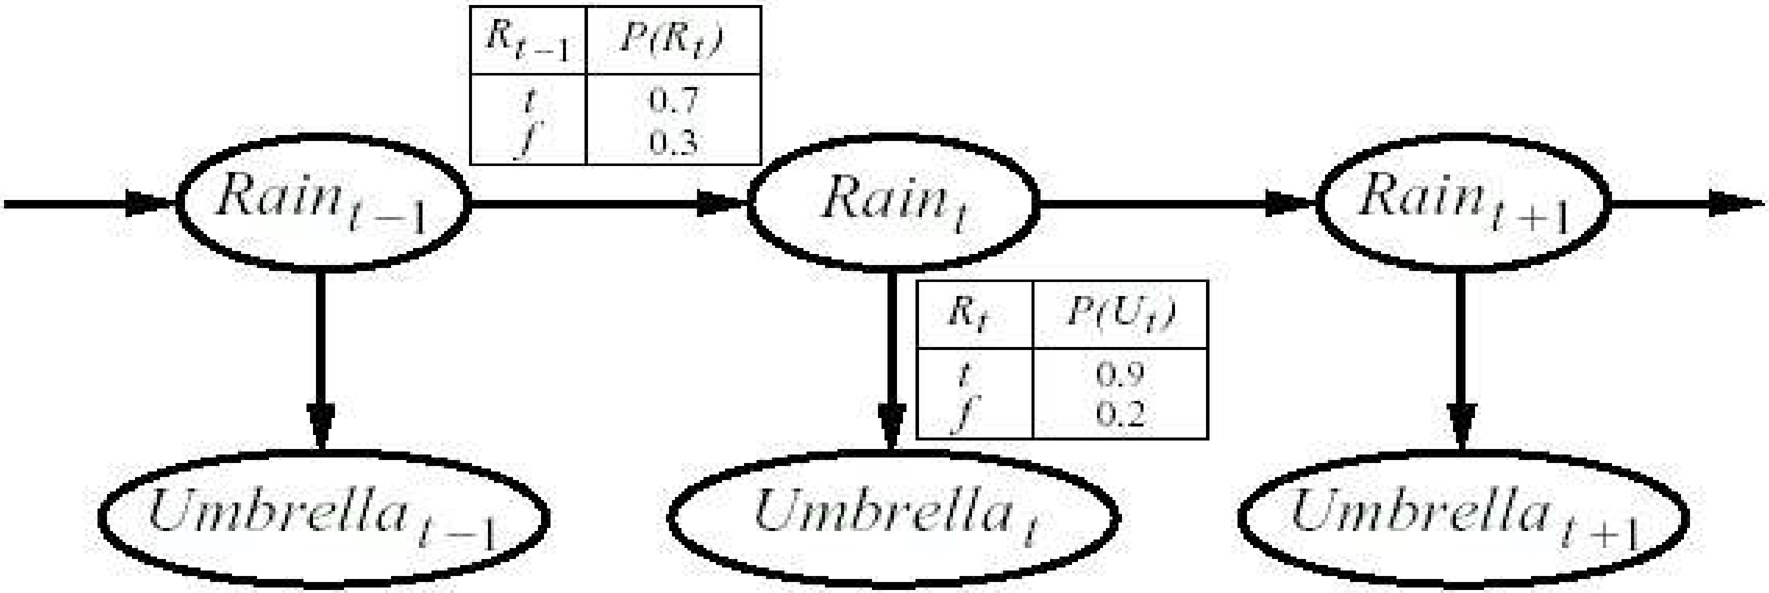

First-order approximation not always true in real world

Possible fixes:

1) Increase order of Markov Process

2) Augment state e.g. add $Temp_t$, $Pressure_t$

Example : robot motion

Augment position and velocity with battery status

Inference Tasks

- Filtering: $P(X_t|e_{1:t})$
  - Belief state: input to decision process of agent
- Prediction: $P(X_{t+k}|e_{1:t})$ for $k > 0$
  - Evaluation of possible action sequences, like filter without full evidence
- Smoothing: $P(X_k|e_{1:t})$for $0 <= k <= t$
  - Better estimate of past states, essential for learning
- Most likely explanation: $\mathrm{argmax}_{x1:t} P(x_{1:t}|e_{1:t})$
  - Speech recognition, decoding over a noisy channel

Filtering

- Aim: Device recursive state estimation:
  - $P(X_{t+1}|e_{1:t+1}) = f(e_{t+1},P(X_t|e_{1:t}))$
  - Current state projected forward from t to t+1, then updated using new evidence e_{t+1}
  - $P(X_{t+1}|e_{1:t+1}) = P(X_{t+1}|e_{1:t},e_{t+1})$
    $= a\,P(e_{t+1}|X_{t+1},e_{1:t})\,P(X_{t+1}|e_{1:t})$ (Bayes rule)
    $= a\,P(e_{t+1}|X_{t+1})P(X_{t+1}|e_{1:t})$ (Markov evidence)
  - Second term: one step prediction of next state
  - First term: updates this with new evidence – directly obtainable from sensor model

Filtering

- $P(X_{t+1}|e_{1:t+1}) = a\, P(e_{t+1}|X_{t+1}) \sum_{xt} P(X_{t+1}|x_t,e_{1:t})P(x_t|e_{1:t})$
  $= a\, P(e_{t+1}|X_{t+1}) \sum_{xt} P(X_{t+1}|x_t)P(x_t|e_{1:t})$ (Markov)
- Within summation first factor is simply the transition model, second is the current state distribution - recursive
- Forward “message” propagation