CSE 591: Human-aware Robotics

Instructor: Dr. Yu ("Tony") Zhang

Location & Times: CAVC 359, Tue/Thu, 9:00--10:15 AM
Office Hours: BYENG 558, Tue/Thu, 10:30--11:30AM

Sep 22, 2016

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

Slides adapted from Henry Kautz and Oregon State University

Plan/Goal Recognition

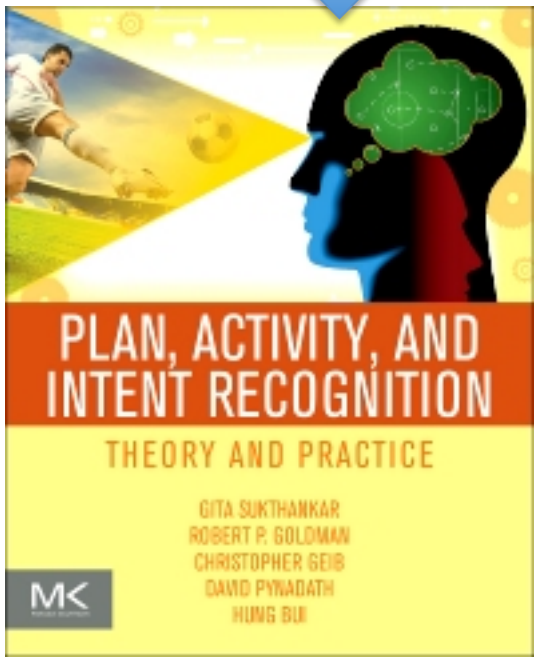

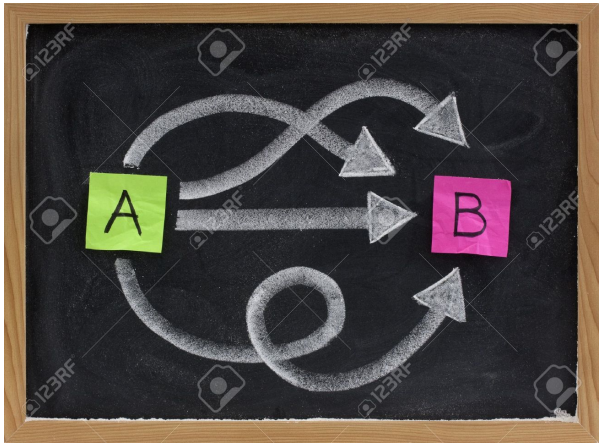

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

Approaches to goal/plan recognition

- Consistency-based
  - Hypothesize & revise
  - Closed-world reasoning
  - Version spaces

- Probabilistic
  - Stochastic grammars
  - Pending sets
  - Layered hidden Markov models
  - Policy recognition
  - Hierarchical hidden semi-Markov models
  - Dynamic probabilistic relational models
  - Dynamic Bayes networks
  - Example application: Assisted Cognition

Can be complementary..
First pick the consistent plans, and check which of them is most likely
(tricky if the agent can make errors)

Hypothesize & Revise

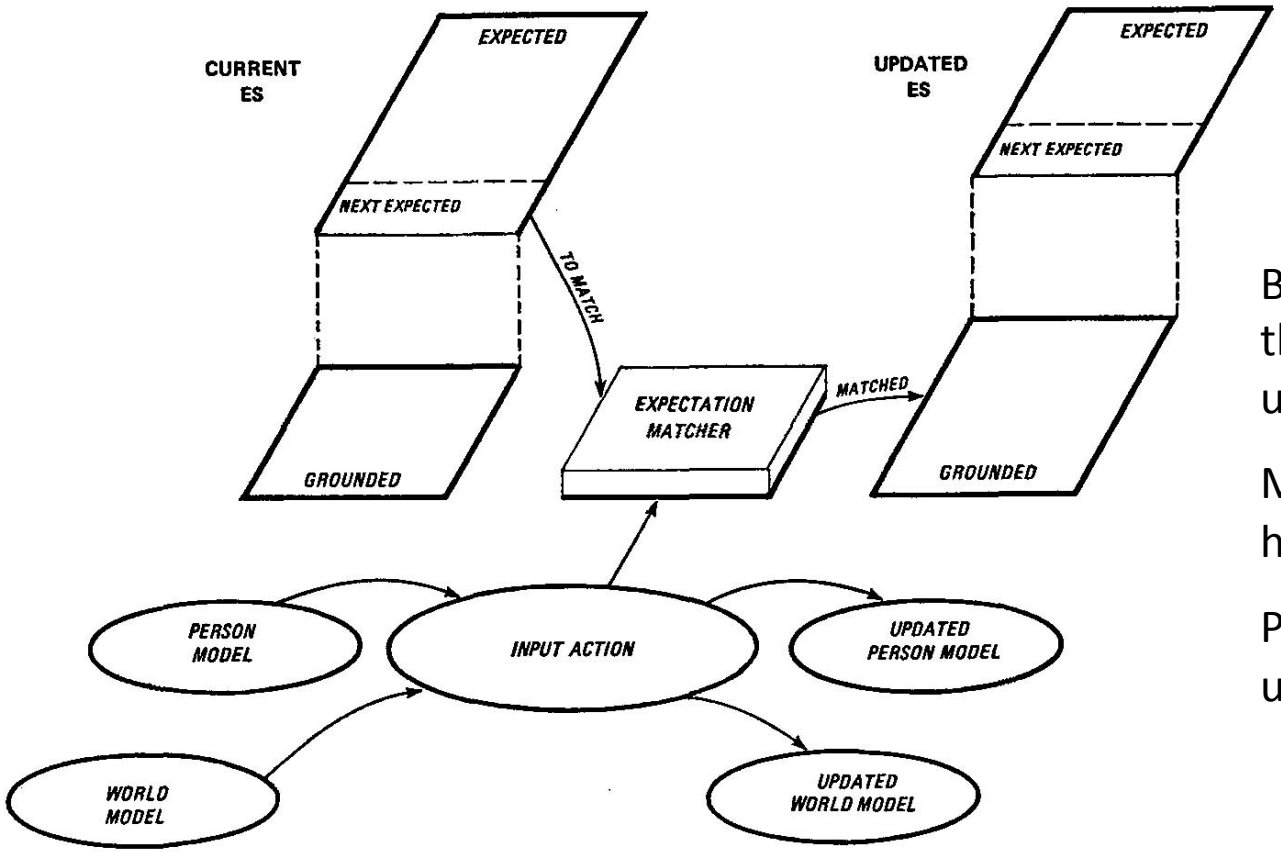

Based on psychological theories of human narrative understanding

Mention of objects suggest hypothesis

Pursue single hypothesis until matching fails

- *The Plan Recognition Problem* C. Schmidt, 1978

Closed-world reasoning

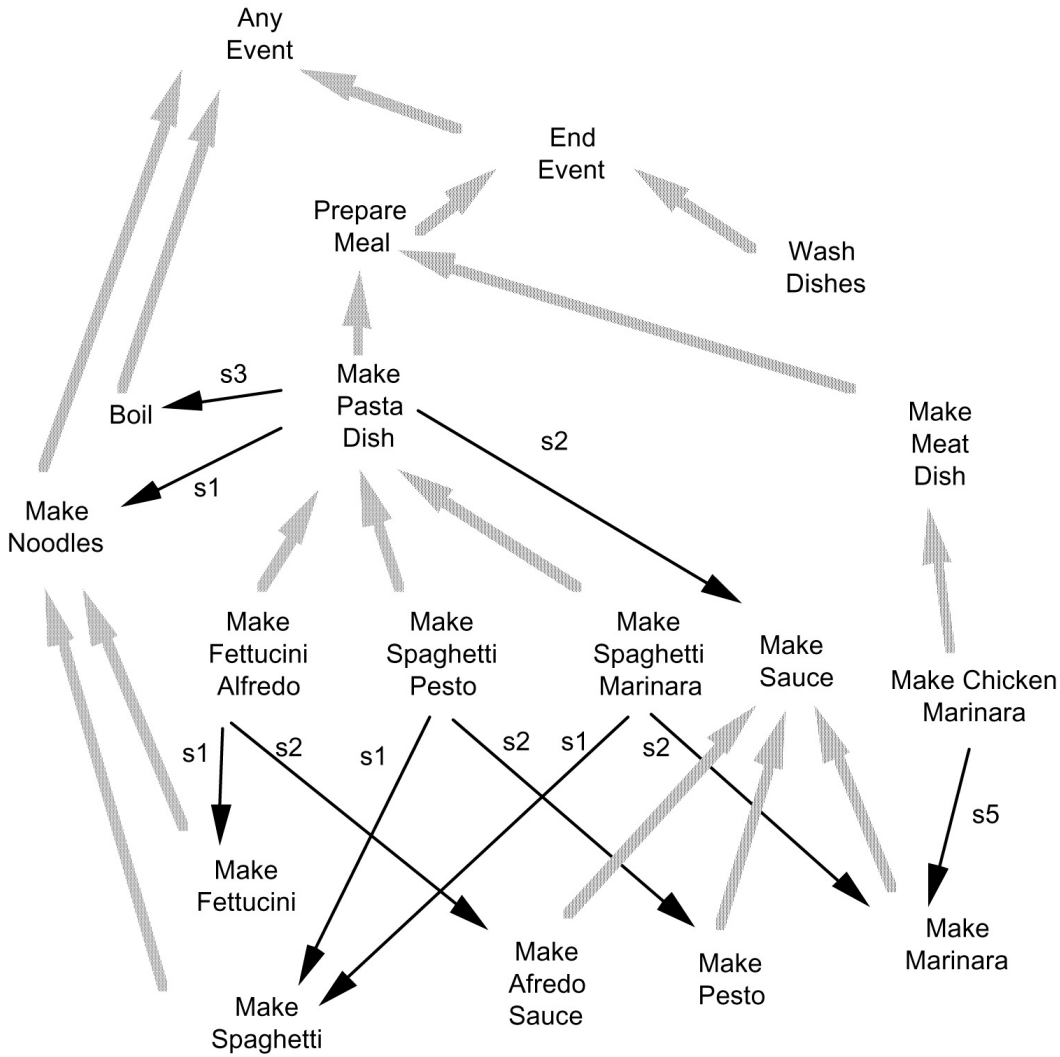

- Infers the minimum set(s) of independent plans that entail the observations
- Observations may be incomplete
- Infallible agent
- Complete plan library

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

- A Formal Theory of Plan Recognition and its Implementation Henry Kautz, 1991

Version Space Algebra

- Recognizes *novel* plans
- Complete observations
- Sensitive to noise
- Start with all goals
- Remove goals when actions are not consistent

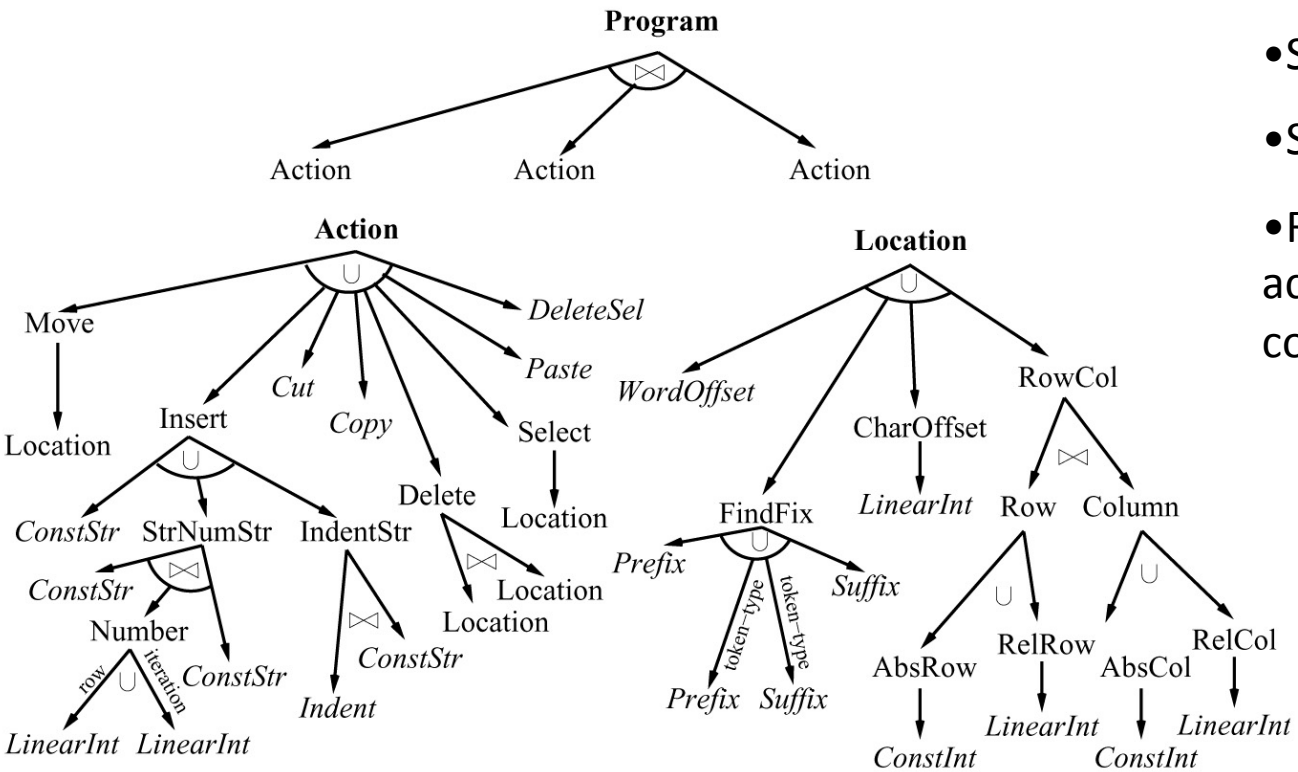

- A sound and fast goal recognizer Lesh & Etzioni
- Programming by Demonstration Using Version Space Algebra Lau, Wolfman, Domingos, Weld.

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

Approaches to goal/plan recognition

- Consistency-based
  - Hypothesize & revise
  - Closed-world reasoning
  - Version spaces

- Probabilistic
  - Stochastic grammars
  - Pending sets
  - Layered hidden Markov models
  - Policy recognition
  - Hierarchical hidden semi-Markov models
  - Dynamic probabilistic relational models
  - Dynamic Bayes networks
  - Example application: Assisted Cognition

Can be complementary..
First pick the consistent plans, and check which of them is most likely
(tricky if the agent can make errors)

Stochastic grammars

Production Rules			Description
S	$\rightarrow AB$	[1.0]	Blackjack → "play game" "determine winner"
A	$\rightarrow CD$	[1.0]	play game → "setup game" "implement strategy"
B	$\rightarrow EF$	[1.0]	determine winner → "eval. strategy" "cleanup"
C	$\rightarrow HI$	[1.0]	setup game → "place bets" "deal card pairs"
D	$\rightarrow GK$	[1.0]	implement strategy → "player strategy"
E	$\rightarrow LKM$	[0.6]	eval. strategy → "dealer down-card" "dealer hits" "player down-card"
	$\rightarrow LM$	[0.4]	eval. strategy → "dealer down-card" "player down-card"
F	$\rightarrow NO$	[0.5]	cleanup → "settle bet" "recover card"
	$\rightarrow ON$	[0.5]	→ "recover card" "settle bet"
G	$\rightarrow J$	[0.8]	player strategy → "Basic Strategy"
	$\rightarrow Hf$	[0.1]	→ "Splitting Pair"
	$\rightarrow bfffH$	[0.1]	→ "Doubling Down"
H	$\rightarrow l$	[0.5]	place bets
	$\rightarrow lH$	[0.5]	
I	$\rightarrow ffI$	[0.5]	deal card pairs
	$\rightarrow ee$	[0.5]	
J	$\rightarrow f$	[0.8]	Basic strategy
	$\rightarrow fJ$	[0.2]	
K	$\rightarrow e$	[0.6]	house hits
	$\rightarrow eK$	[0.4]	
L	$\rightarrow ae$	[1.0]	Dealer downcard
M	$\rightarrow dh$	[1.0]	Player downcard
N	$\rightarrow k$	[0.16]	settle bet
	$\rightarrow kN$	[0.16]	
	$\rightarrow j$	[0.16]	
	$\rightarrow jN$	[0.16]	
	$\rightarrow i$	[0.18]	
	$\rightarrow iN$	[0.18]	
O	$\rightarrow a$	[0.25]	recover card
	$\rightarrow aO$	[0.25]	
	$\rightarrow b$	[0.25]	
	$\rightarrow bO$	[0.25]	

Symbol	Domain-Specific Events (Terminals)
a	dealer removed card from house
b	dealer removed card from player
c	player removed card from house
d	player removed card from player
e	dealer added card to house
f	dealer dealt card to player
g	player added card to house
h	player added card to player
i	dealer removed chip
j	player removed chip
k	dealer pays player chip
l	player bets chip

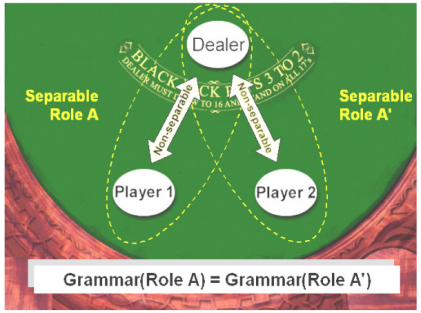

Seperable, Independent roles

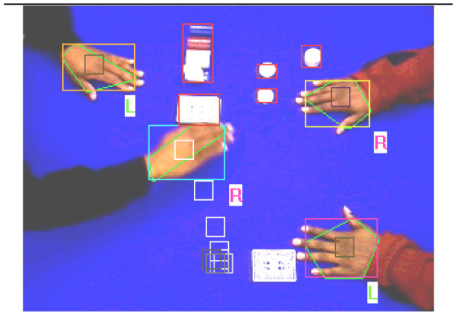

Vision system tracking the game.

CF grammar w/ probabilistic rules

Successful for highly structured tasks (e.g. playing cards)

Problems: errors, context

- Huber, Durfee, & Wellman, "The Automated Mapping of Plans for Plan Recognition", 1994
- Darnell Moore and Irfan Essa, "Recognizing Multitasked Activities from Video using Stochastic Context-Free Grammar", AAAI-02, 2002.

Pending sets

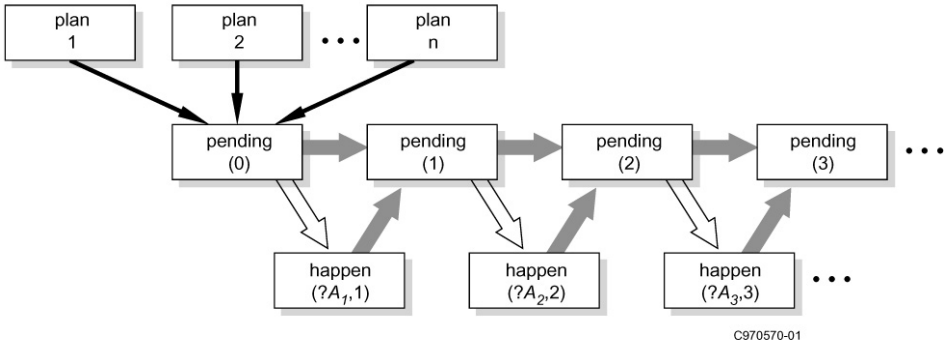

Explicitly models the agent's "plan agenda" using Poole's "probabilistic Horn abduction" rules

Handles multiple concurrent interleaved plans & negative evidence

Number of different possible pending sets can grow exponentially

Context problematic?
Metric time?

Happen(X,T+1) ←
  Pending(P,T),
  X in P,
  Pick(X,P,T+1).

Pending(P',T+1)←
  Pending(P,T),
  Leaves(L),
  Progress(L, P, P', T+1).

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

- A new model of plan recognition. Goldman, Geib, and Miller
- Probabilistic plan recognition for hostile agents. Geib, Goldman

Layered hidden Markov models

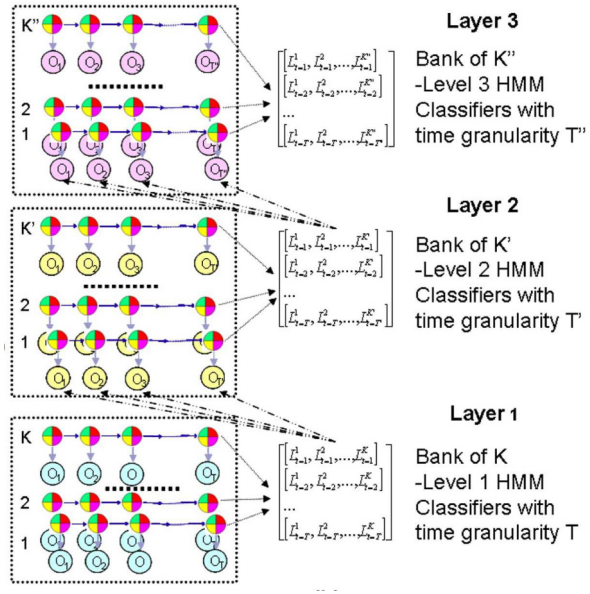

Cascade of HMM's, operating at different temporal granularities

Inferential output at layer K is "evidence" for layer K +1

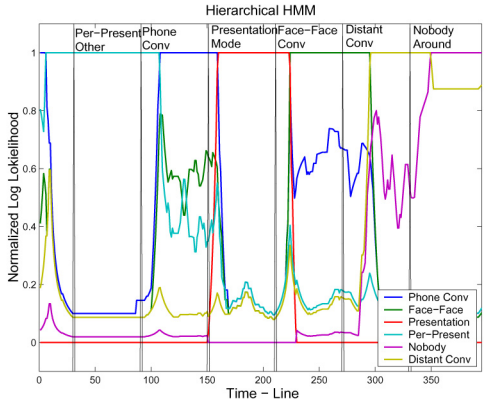

- N. Oliver, E. Horvitz, and A. Garg. Layered Representations for Recognizing Office Activity, *Proceedings of the Fourth IEEE International Conference on Multimodal Interaction (ICMI 2002)*

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

Policy Recognition

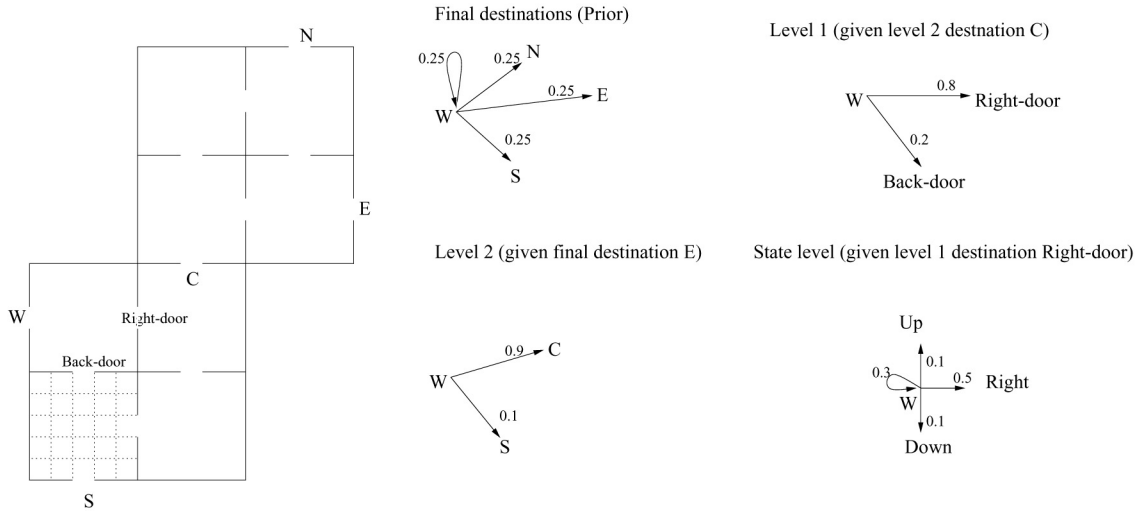

(a) The environment

(b) Parameters of the AHMM

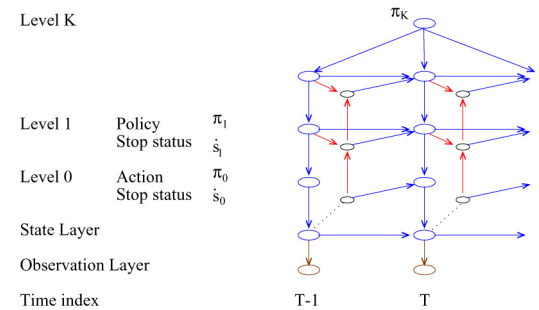

Model agent using hierarchy of abstract policies (e.g. abstract by spatial decomposition)

Compute the conditional probability of top-level policy given observations

Compiled into DBN

- Tracking and Surveillance in Wide-Area Spatial Environments Using the Hidden Markov Model. Hung H. Bui, Svetha Venkatesh and West.
- Bui, H. H., Venkatesh, S., and West, G. (2000) On the recognition of abstract Markov policies. Seventeenth National Conference on Artificial Intelligence (AAAI-2000), Austin, Texas

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

Hierarchical hidden semi-Markov models

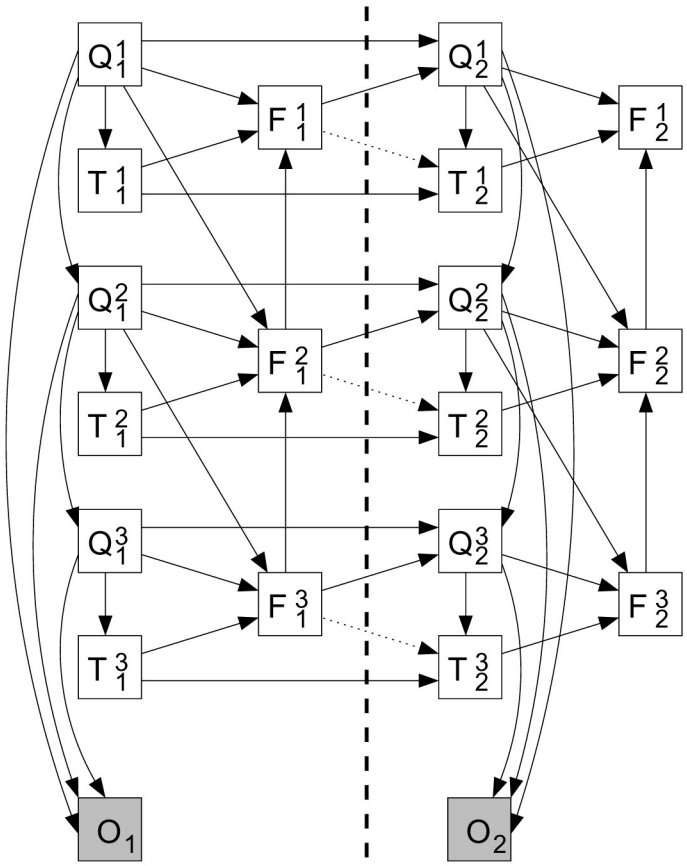

Combine hierarchy (function call semantics) with metric time

Compile to DBN

Time nodes represent a distribution over the time of the next state "switch"

"Linear time" smoothing

– Research issues – parametric time nodes, varying granularity

- Hidden semi-Markov models (segment models) Kevin Murphy. November 2002.
- HSSM: Theory into Practice, Deibel & Kautz, forthcoming.

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

Dynamic probabilistic relational models

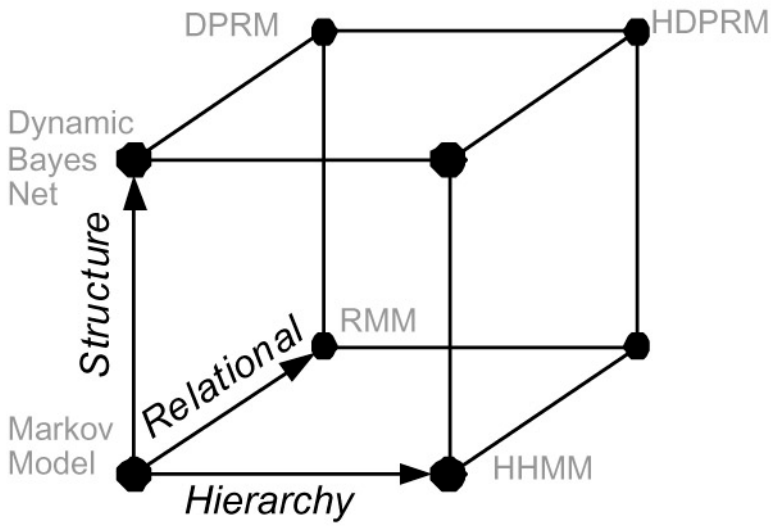

PRM - reasons about classes of objects and relations

Lattice of classes can capture plan abstraction

DPRM – efficient approximate inference by Rao-Blackwellized particle filtering

Open: approximate smoothing?

- Friedman, N., L. Getoor, D. Koller, A. Pfeffer. Learning Probabilistic Relational Models. *IJCAI-99*, Stockholm, Sweden (July 1999).
- Relational Markov Models and their Application to Adaptive Web Navigation, Anderson, Domingos, Weld 2002.
- Dynamic probabilistic relational models, Anderson, Domingos, Weld, forthcoming.

Dynamic Bayesian Network

Time and Change in Probabilistic Reasoning

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

Markov processes (Markov chains)

Construct a Bayes net from these variables: parents?

Markov assumption: $\mathbf{X}_t$ depends on **bounded** subset of $\mathbf{X}_{0:t-1}$

First-order Markov process: $\mathbf{P}(\mathbf{X}_t|\mathbf{X}_{0:t-1}) = \mathbf{P}(\mathbf{X}_t|\mathbf{X}_{t-1})$
Second-order Markov process: $\mathbf{P}(\mathbf{X}_t|\mathbf{X}_{0:t-1}) = \mathbf{P}(\mathbf{X}_t|\mathbf{X}_{t-2}, \mathbf{X}_{t-1})$

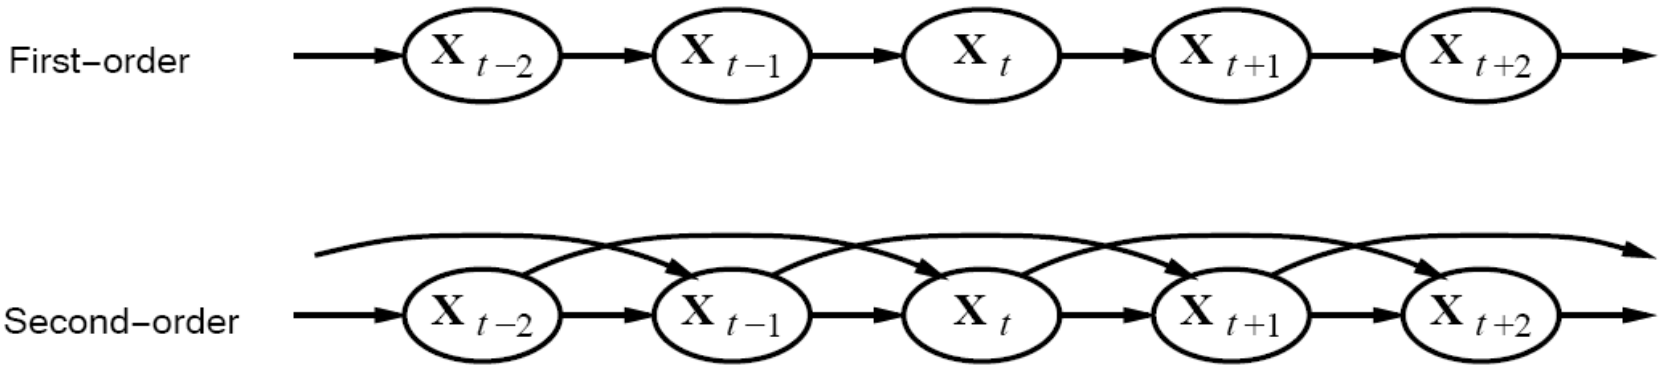

Sensor Markov assumption: $\mathbf{P}(\mathbf{E}_t|\mathbf{X}_{0:t}, \mathbf{E}_{0:t-1}) = \mathbf{P}(\mathbf{E}_t|\mathbf{X}_t)$

Stationary process: transition model $\mathbf{P}(\mathbf{X}_t|\mathbf{X}_{t-1})$ and sensor model $\mathbf{P}(\mathbf{E}_t|\mathbf{X}_t)$ fixed for all t

Time and uncertainty

The world changes; we need to track and predict it

Diabetes management vs vehicle diagnosis

Basic idea: copy state and evidence variables for each time step

$\mathbf{X}_t$ = set of unobservable state variables at time t
e.g., $BloodSugar_t$, $StomachContents_t$, etc.

$\mathbf{E}_t$ = set of observable evidence variables at time t
e.g., $MeasuredBloodSugar_t$, $PulseRate_t$, $FoodEaten_t$

This assumes **discrete time**; step size depends on problem

Notation: $\mathbf{X}_{a:b} = \mathbf{X}_a, \mathbf{X}_{a+1}, \ldots, \mathbf{X}_{b-1}, \mathbf{X}_b$

Example

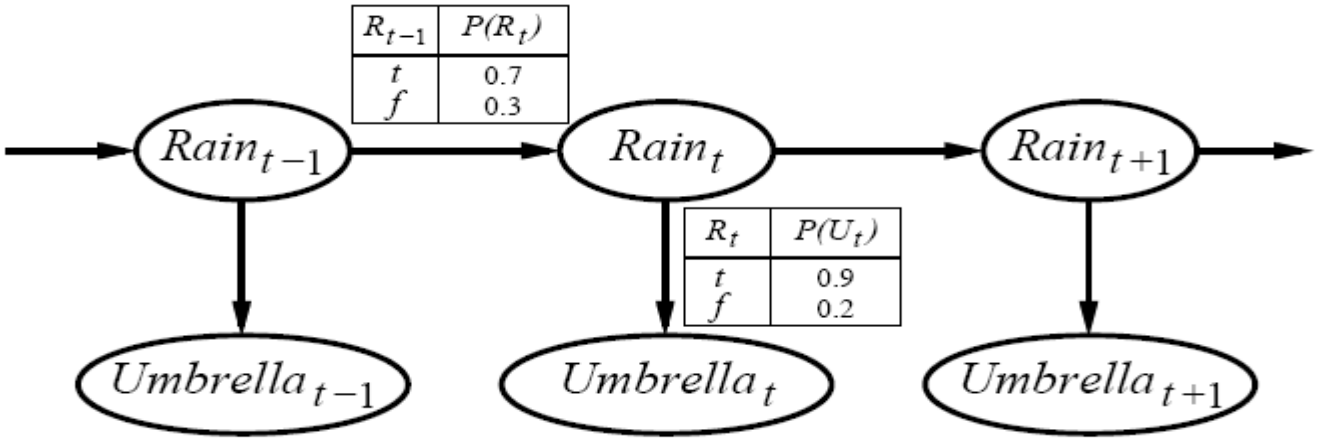

Issues: 1-order Markov; huge CPT table

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 $Battery_t$

Dynamic Bayesian networks

$\mathbf{X}_t$, $\mathbf{E}_t$ contain arbitrarily many variables in a replicated Bayes net

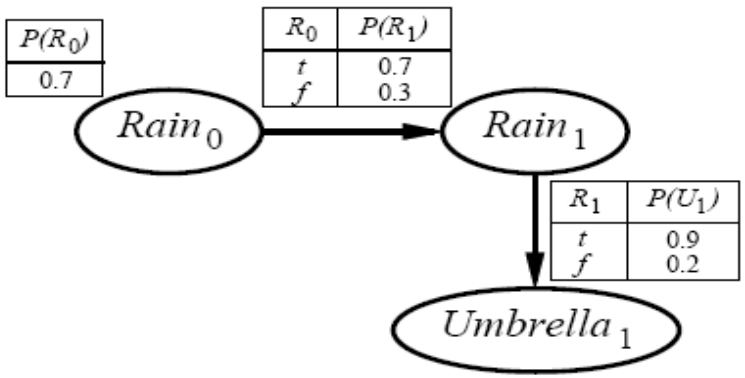

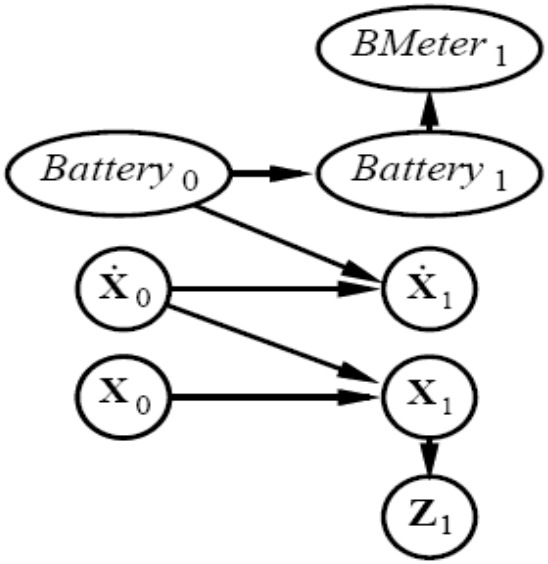

Dynamic Bayes Networks are "templates" for specifying the relation between

the values of a random variable across time-slices

→e.g. How is Rain at time t related to Rain at time t+1?

We call them templates because they need to be expanded (unfolded) to the

required number of time steps to reason about the connection between

variables at different time points

Special Cases of DBNs are well known in the literature

- Restrict number of variables per state
  - Markov Chain: DBN with one variable that is fully observable
  - Hidden Markov Model: DBN with only one state variable that is hidden and can be estimated through evidence variable(s)
- Restrict the type of CPD
  - Kalman Filters: DBN where the system transition function as well as the observation variable are *linear gaussian*
    - The advantage of Gaussians is that the posterior distribution remains Gaussian

Kalman filters

Modelling systems described by a set of continuous variables,
e.g., tracking a bird flying—$\mathbf{X}_t = X, Y, Z, \dot{X}, \dot{Y}, \dot{Z}$.
Airplanes, robots, ecosystems, economies, chemical plants, planets, ...

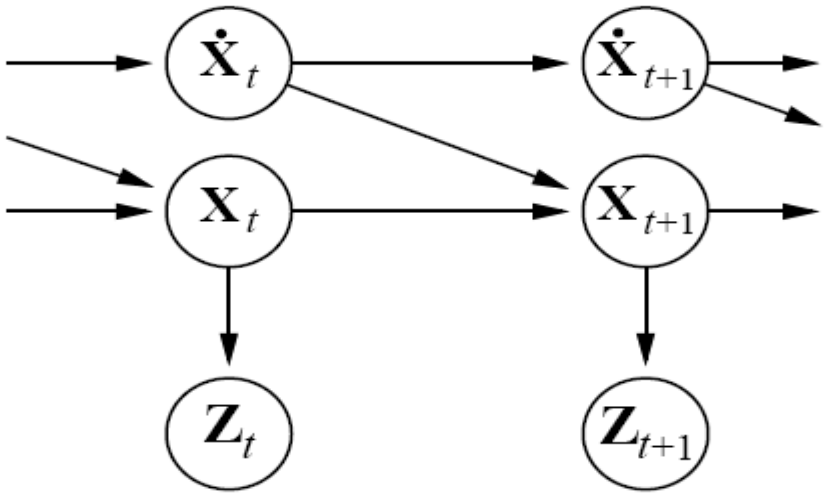

Gaussian prior, linear Gaussian transition model and sensor model

DBNs vs Kalman filters

Every Kalman filter model is a DBN, but few DBNs are KFs;
real world requires non-Gaussian posteriors

E.g., where are bin Laden and my keys? What's the battery charge?

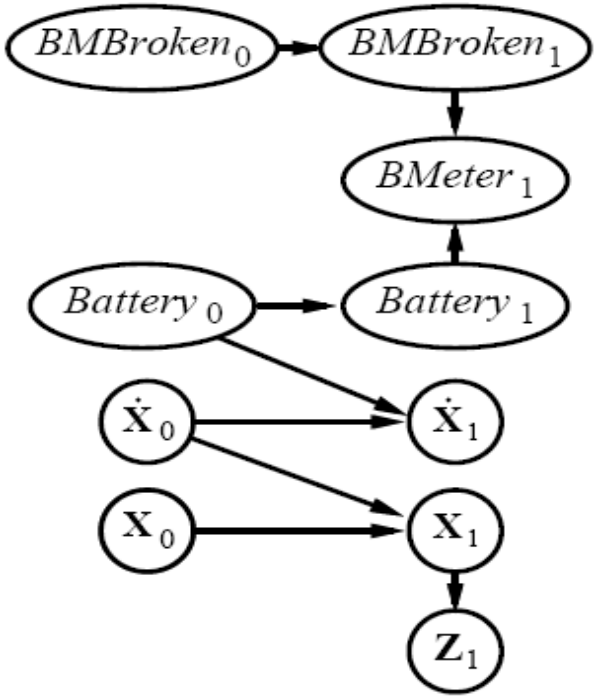

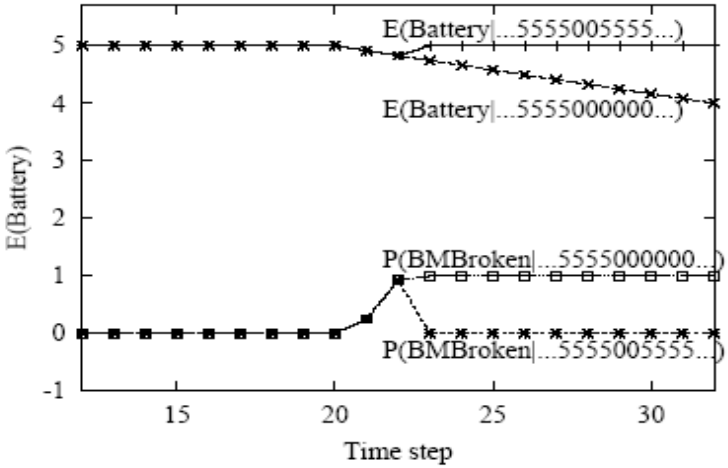

Plan Recognition Approaches based on setting up DBNs

Dynamic Bayes nets

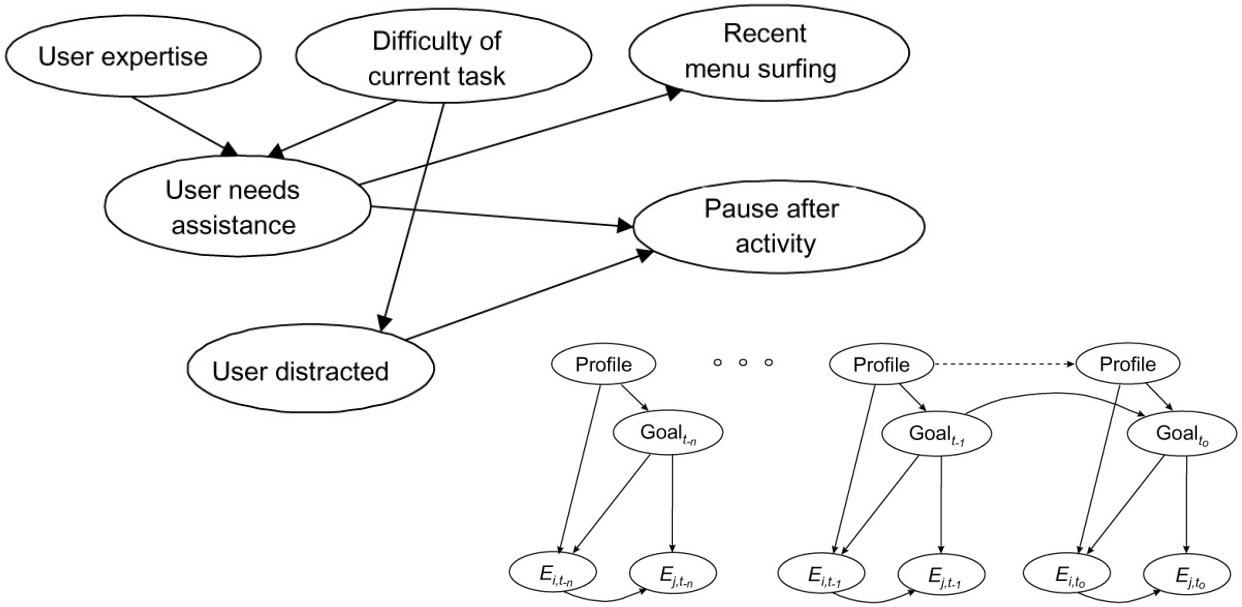

Models relationship between user's recent actions and goals (help needs)

Probabilistic goal persistence

- E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and K. Rommelse. The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users. *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, July 1998.
- Towards a Bayesian model for keyhole plan recognition in large domains Albrecht, Zukermann, Nicholson, Bud

Dynamic Bayesian Nets

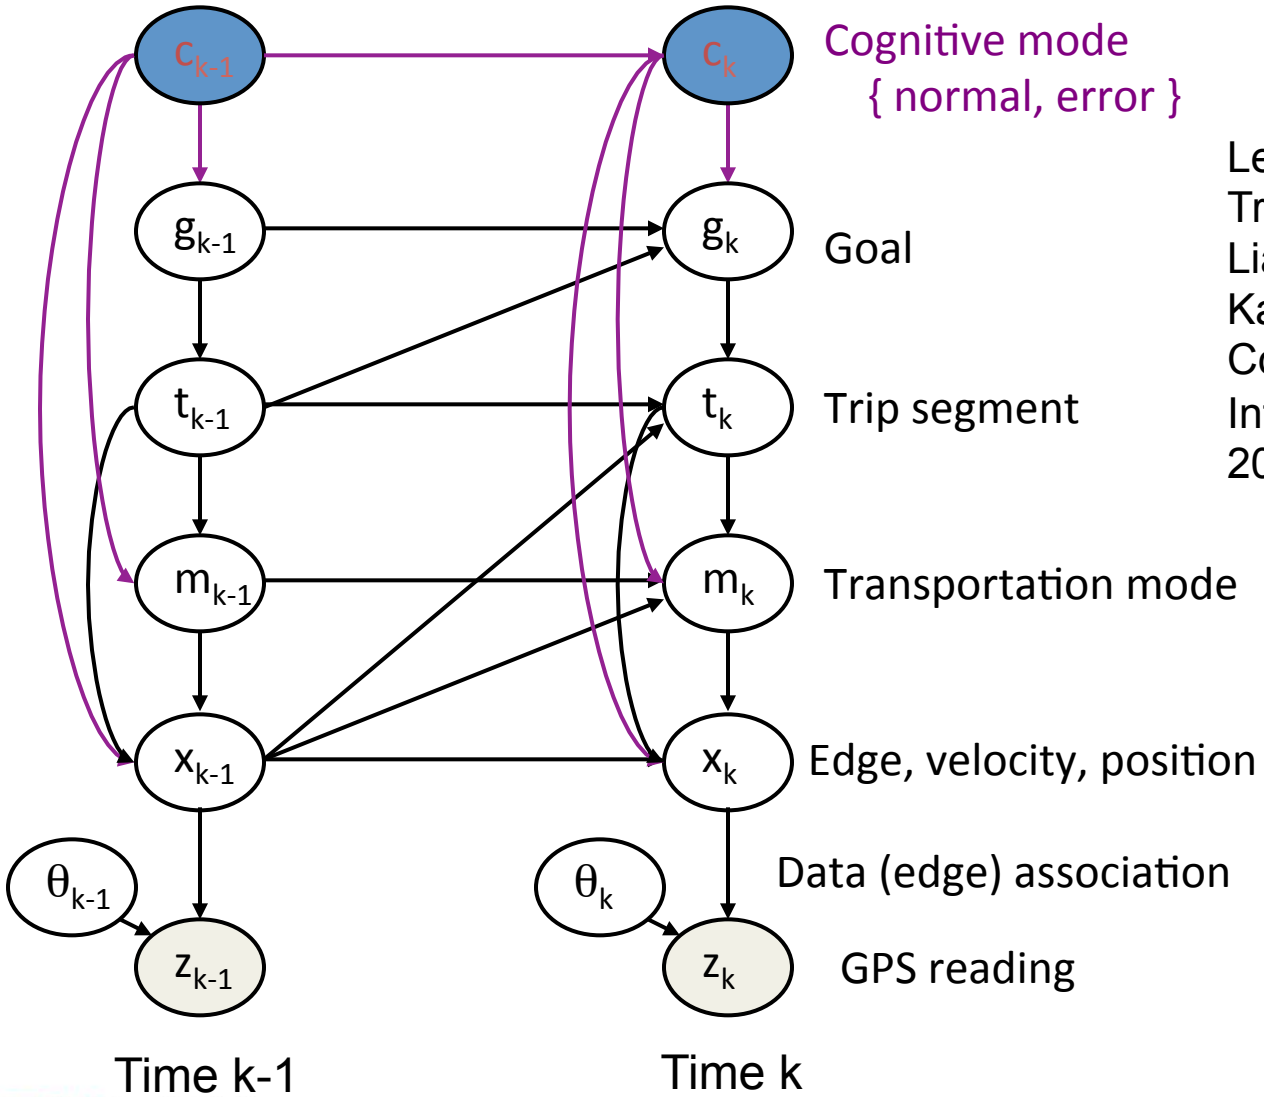

Learning and Inferring Transportation Routines Lin Liao, Dieter Fox, and Henry Kautz, Nineteenth National Conference on Artificial Intelligence, San Jose, CA, 2004.

Applications

Assisted cognition

Computer systems that improve the independence and safety of people suffering from cognitive limitations by…

- Understanding human behavior from low-level sensory data
  - Using commonsense knowledge
  - Learning individual user models
- Actively offering prompts and other forms of help as needed
- Alerting human caregivers when necessary

http://www.cs.washington.edu/assistcog/

Activity Compass

- Zero-configuration personal guidance system
  - Learns model of user's travel on foot, by public transit, by bike, by car
  - Predicts user's next destination, offers proactive help if lost or late
- Integrates user data with external constraints
  - Maps, bus schedules, calendars, …
  - EM approach to clustering & segmenting data

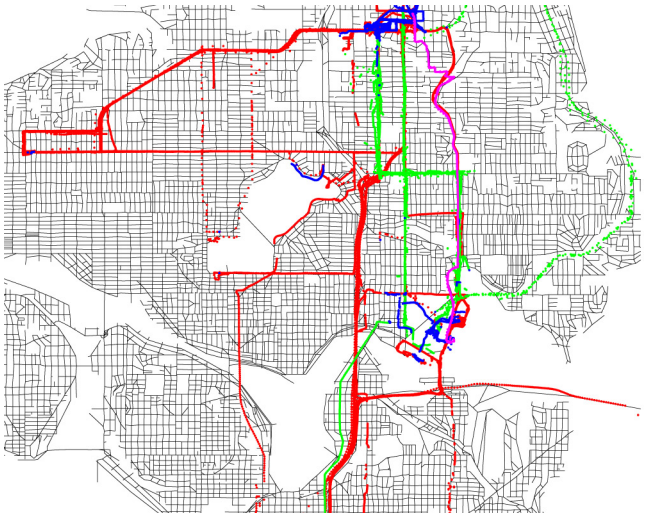

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

The Activity Compass Don Patterson, Oren Etzioni, and Henry Kautz (2003)

Activity of daily living monitor & prompter

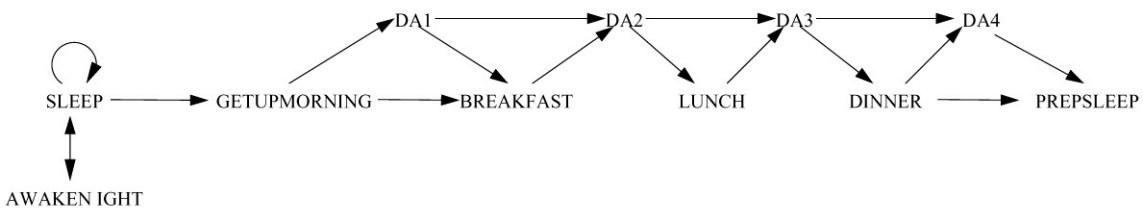

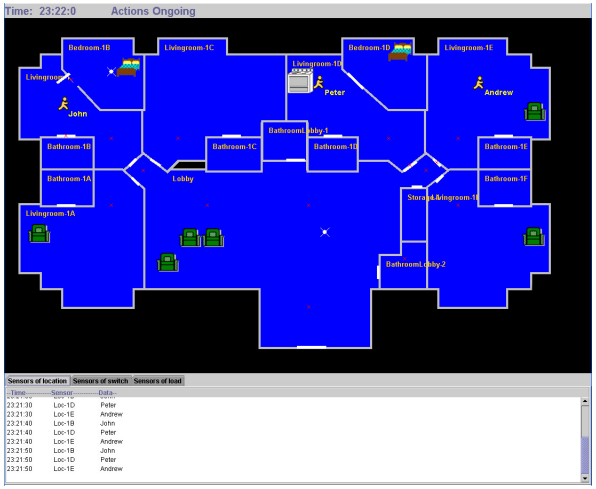

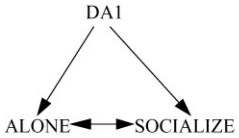

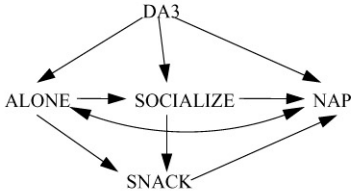

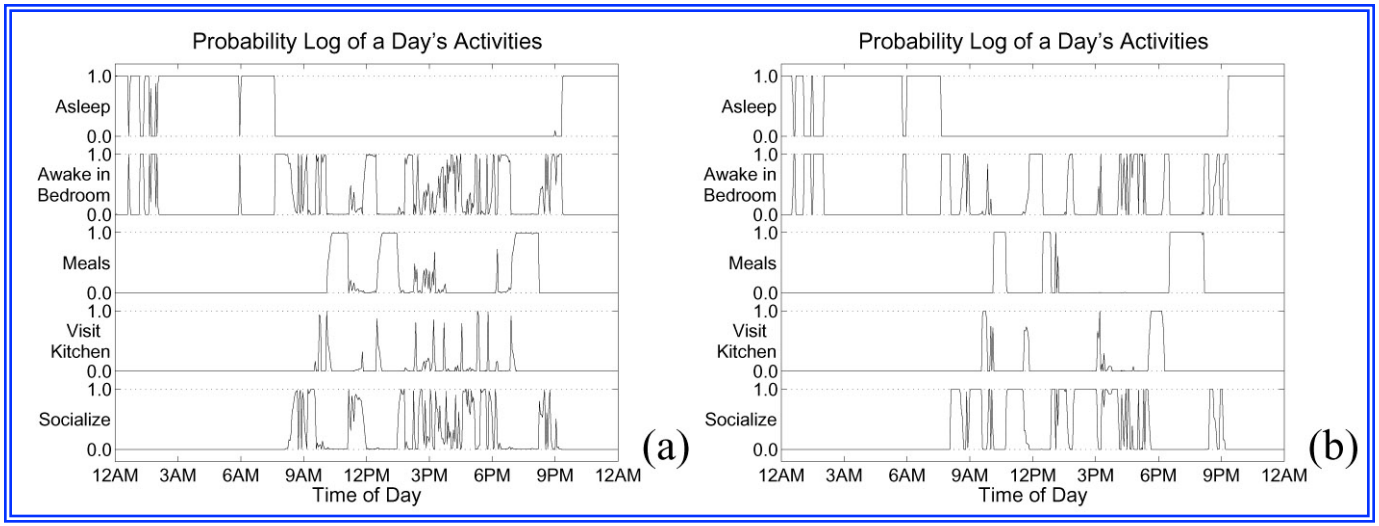

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

Foundations of Assisted Cognition Systems. Kautz, Etzioni, Fox, Weld, and Shastri, 2003

Recognizing unexpected events using online model selection

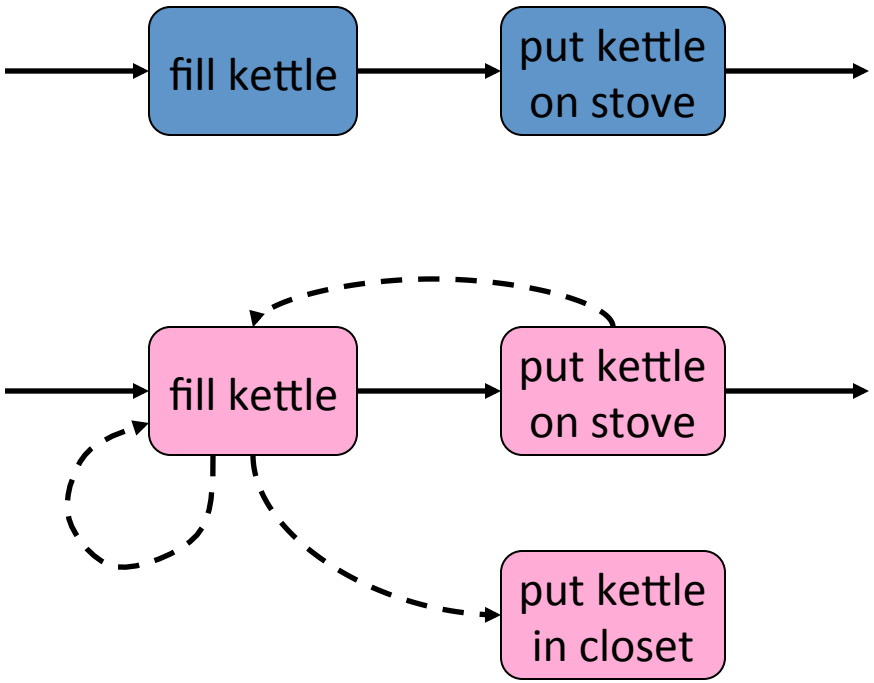

- User errors, abnormal behavior
- Select model that maximizes likelihood of data:
  - Generic model
  - User-specific model
  - Corrupt (impaired) user model
- Neurologically-plausible corruptions
  - Repetition
  - Substitution
  - Stalling

Ira A. Fulton Schools of Engineering
ARIZONA STATE UNIVERSITY

Fox, Kautz, & Shastri