CAP6671 Intelligent Systems

Lecture 20:

Human-Agent Teamwork

Instructor: Dr. Gita Sukthankar
Email: gitars@eecs.ucf.edu

Schedule: T & Th 9:00-10:15am

Location: HEC 302
Office Hours (in HEC 232):
T & Th 10:30am-12

Why have human-agent teams?

Why have human-agent teams?

- Enhance human information processing capabilities
- Augment human reasoning with agents capable of doing task-specific calculations
- Monitor human performance

Research Projects

- Enhancing team search operations (Sukthankar, Sycara, and Giampapa, 2007)
- Cognitive load modeling (Fan and Yen, 2007)

Models of Human-Agent Teams

Agent aiding teamwork activities

Agent acting as team member

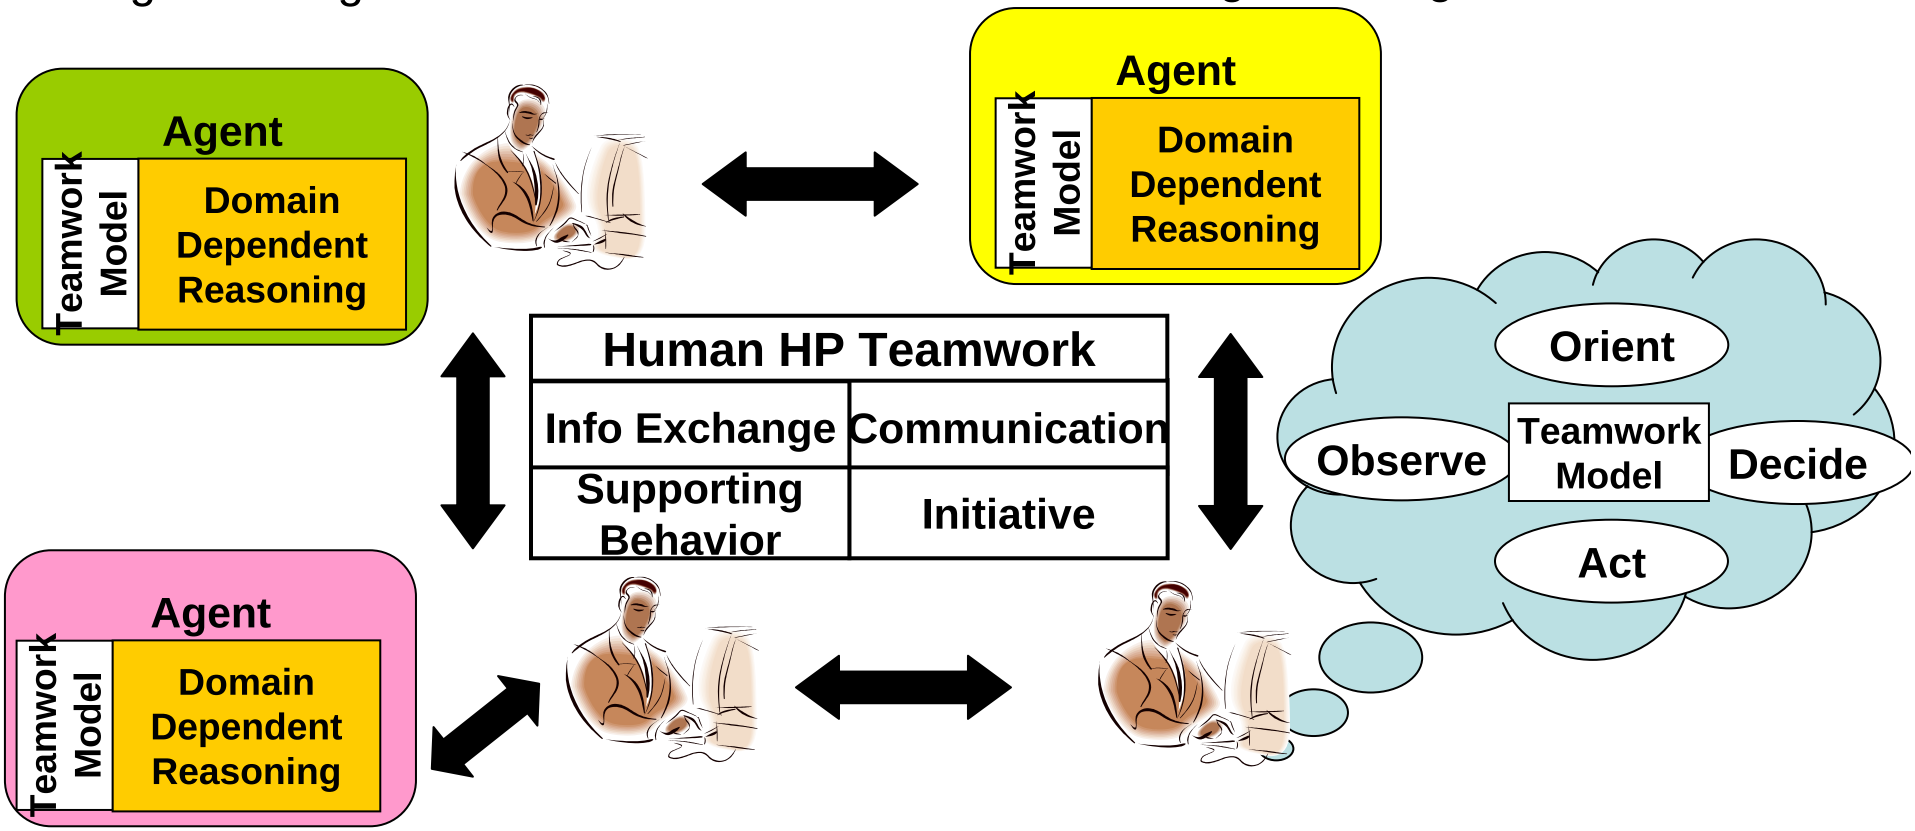

Agent serving as human team member assistant

Principle: By equipping software agents with domain-independent models of teamwork, agents are 1) able to assume the role of full team members 2) more capable of recognizing, monitoring and aiding teamwork activities between human teammates.

Towards Agent-Based Teamwork

- Establish a baseline of human-only team performance
- Understand how people organize, divide tasks, plan, execute, etc.
- Discover opportunities in which humans can be aided
- Understand the types of communication and coordination humans use
- Understand the strategies people use
  - E.g. Search widely and hastily or carefully in a limited areas?
  - Identify cues humans use when changing strategies

Implications of Pilot Experiments

- Understand the types of measures that can be used in agent-aiding experiments
- Future task manipulations
- Future task variations to be tested
- “Points of insertion” of agent aiding
- Utility of different agent aiding schemes
- Logging and data collection requirements

Opportunities for Error in Human Teams: Gaps in Search Patterns

- Influencing factors:
  - Being conditioned by terrain features
    (e.g. ridges, hills, hedge rows, tree lines, etc.)
  - Being conditioned by cultural features
    (e.g. buildings, courtyards, radial vs. grid layout of village roads, etc.)
  - Discovery of sought object outside of search path/area
  - Taking short-cuts during search
  - Difficulty orienting self with terrain, map, other players

Opportunities for Error in Human Teams

- Poor priority assignments in the search plan due to false clues or hunches
  - Strong but unmotivated convictions about where to locate sought objects
  - Duplicated, unnecessary and repeated searching of “preferred” areas
- Poor pacing during time-critical tasks
  - Initial difficulty estimating the pace at which to perform search
  - Difficulty in maintaining pace until the final critical minutes

Experimental setting

- 17 teams of 3 subjects each
- Goal:
  - To find & crush as many bottles as possible
  - Bottle placement follows a different probability distribution each session
  - Individual and team score: 1 point per bottle
- Subjects are unconstrained to:
  - Organize their team, decide and allocate roles,
  - Planning before/during execution, execute plans, etc.
- Three sessions:
  - One practice session to familiarize subjects with the software environments and available tools (e.g. 2-D map, binoculars, compass, watch, commands)
  - Two test conditions:
    - Total bottle count is revealed to the team ("***<u>Known</u>***" condition)
    - Total bottle count unknown ("***<u>Unknown</u>***" condition)
    - The two conditions were ran in alternating sequences for the subject teams
- Subject audio communications
  - Were recorded during all three sessions
  - Audio logs are being manually transcribed and coded

Experimental Testbed: 2-D Map + Instruments

- Implementation in ***Operation Flashpoint***
- Subjects use virtual tools such as compass, binoculars, and watch
- Communication takes place over headsets and is recorded for post-experiment analysis using ***TeamSpeak***

Subject Viewpoint

- Scenario requires subjects to locate and crush bottles that are hidden on the map
- Screenshot shows a practice area with a high density of bottles

Bottle Collection Interface

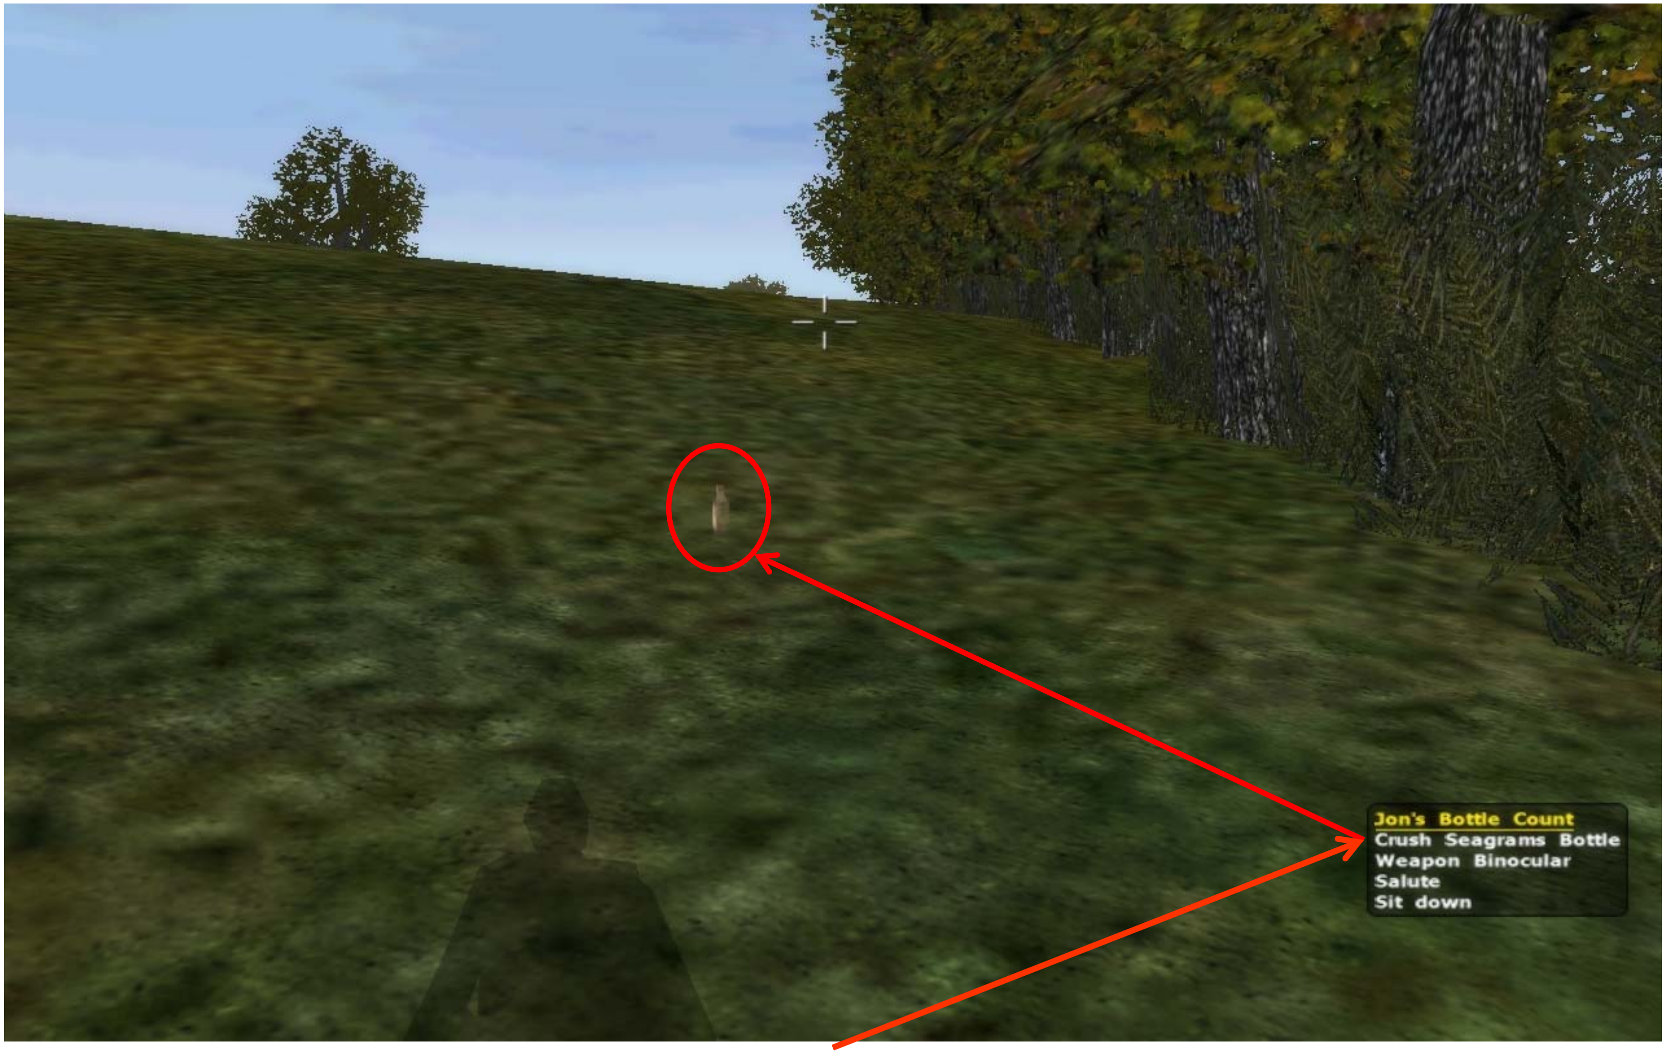

The command to “crush” a bottle only becomes available when the player is near it.

Communication Categories

- Increasing team situational awareness
  - Communicating teammate locations

    Where are you? I am here.
  - Terrain/map features

    There are some impassable hedges by the NATO headquarters.
  - Bottle locations

    I see a bottle by the British flag.
- Hints, sharing individual search strategies

  How do I scan for bottles?
- Team Planning, Before or During Execution
  - Role allocation

    The character with the farthest 2-D sensing should scan the most map squares.
  - Division of execution space

    You scan the village.

Communication Categories

- Task Progress Monitoring
  - Bottle counts
    We have 20 bottles and need 15 more in 10 minutes.
  - Coverage progress
    I have covered this sector.
- Internal Beliefs
  - Object density
    I think that there is a hidden cache of bottles.
  - Object locations
    There seem to be more bottles in the town.
- Miscellaneous
  - System problems
    My keyboard locked up!
  - Personal discussions
    All these bottles lying around … looks like our apartment, eh?

Communication Patterns

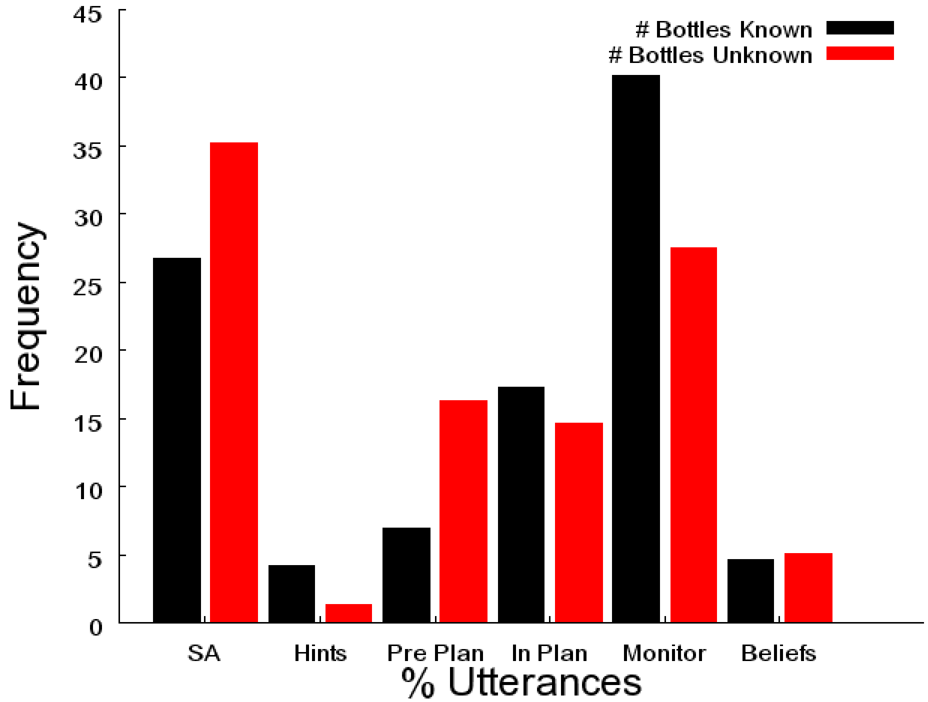

- Communication frequency averaged across four teams of subjects
- ***# Bottles Known*** condition
  - Subjects appear to spend more time monitoring their task progress
- ***# Bottles Unknown*** condition
  - Subjects appear to spend more time planning prior to execution
  - More communication devoted to increasing team's situation awareness

Search Pattern Coordination Analysis

- Subjects indicate their search paths on a paper map.
- Patterns are used to estimate subject and team area coverage.
- Drawbacks:
  - Coarse-grained encoding of the search space (24 grid cells)
  - Does not illustrate ***all*** holes in search patterns due to:
    - Gaps in line of sight caused by variations of terrain elevation
    - Obstructions: buildings, hedges, trees, low brush
  - Self-reported by subjects, usually post experiment session
- Benefits:
  - Quantifies degree of terrain covered and search path overlap
  - Indicates ***some*** holes in team-coordinated search patterns
  - Highlights search area responsibilities that subjects assumed

Terrain Coverage by Worst Performing Team

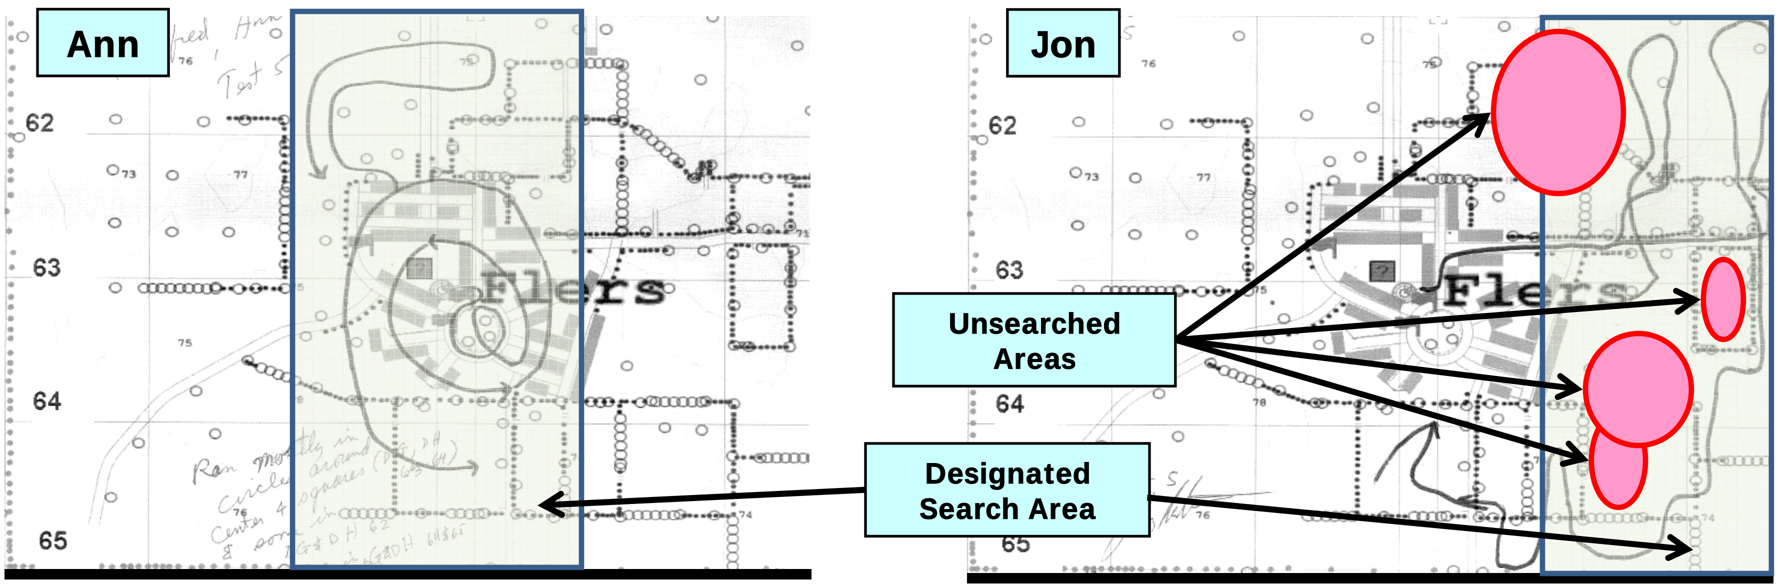

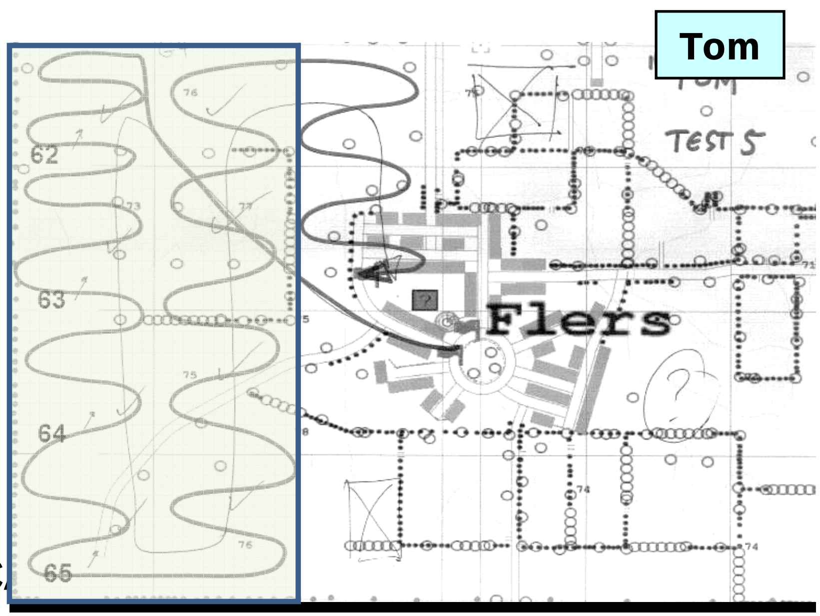

Team 4 – "Bottle Count Unknown" Condition
Ann (Novice), Jon (Novice), Tom (Medium)
(6 + 7 + 2) / 42 = 35.71% bottles found
100% Terrain Coverage (33%, 46%, 50%)
29% Terrain Coverage Overlap

Notwithstanding the holes in Jon's declared search pattern, he found more bottles (7) than Tom (2), who declared a `complete' search pattern.

Observed Egocentric Bias

- Most human subjects preferred to search the village
  - Rationale: *There seem to be more bottles in the village.*
  - Hypothesis:
    - The bottles were ***easier to find*** in in the village.
    - Subjects sought quicker & easier gratification
- Egocentric bias:
  - Leaving mutually-agreed upon search area to search village
  - Other teammates were doing the same
    - Individuals either did not know, or
    - They were not thinking about the consequences
- Negative consequences:
  - Breaking team commitments to roles and strategy, increased need for monitoring
  - Unnecessary duplication of “easy” effort
  - Less time to search areas with more difficult terrain

Conclusion: Where Agents May Help

- Aiding the individual in a team
  - Assist in individual performance monitoring/improvement functions:
    - Completeness of terrain search
    - Compliance with commitments
  - Be a personal coach for maintaining proper execution pace and technique
  - Be an on-demand guide when individual becomes disoriented
- Helping the team
  - Assist in team performance monitoring/improvement functions
    - Collecting, disseminating and appropriately displaying situation awareness information
    - Monitoring team member compliance to committed goals
    - Suggesting corrective actions when team under-performs
  - Proactively retrieve performance history and outcome of similar tasks, and indicate parallels with current task

Cognitive Load Modeling (Fan & Yen)

- Theories of human teamwork say:
  - To perform effectively teams need to establish shared mental models of the situation
  - Establishing shared mental models requires communication
- Problem: Too much communication causes information overload!
- Solution:
  - Agents model human information load using an HMM formalism to recognize hidden state
  - Reduce communication to overloaded agents

Cognitive Load

- Human information-processing system:
  - Sensory store (one second)
  - Working memory (approx 7 elements for several seconds)
  - Long-term memory
- Cognitive load arises due to the limitations of working memory.
- Can be measured:
  - Directly (physiological measures)
  - Indirectly (performance of a secondary task)

CPDs for HMMs

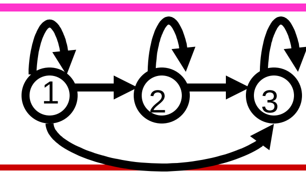

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ 0 & a_{22} & a_{23} \\ 0 & 0 & a_{33} \end{pmatrix}$$

A=state transition matrix

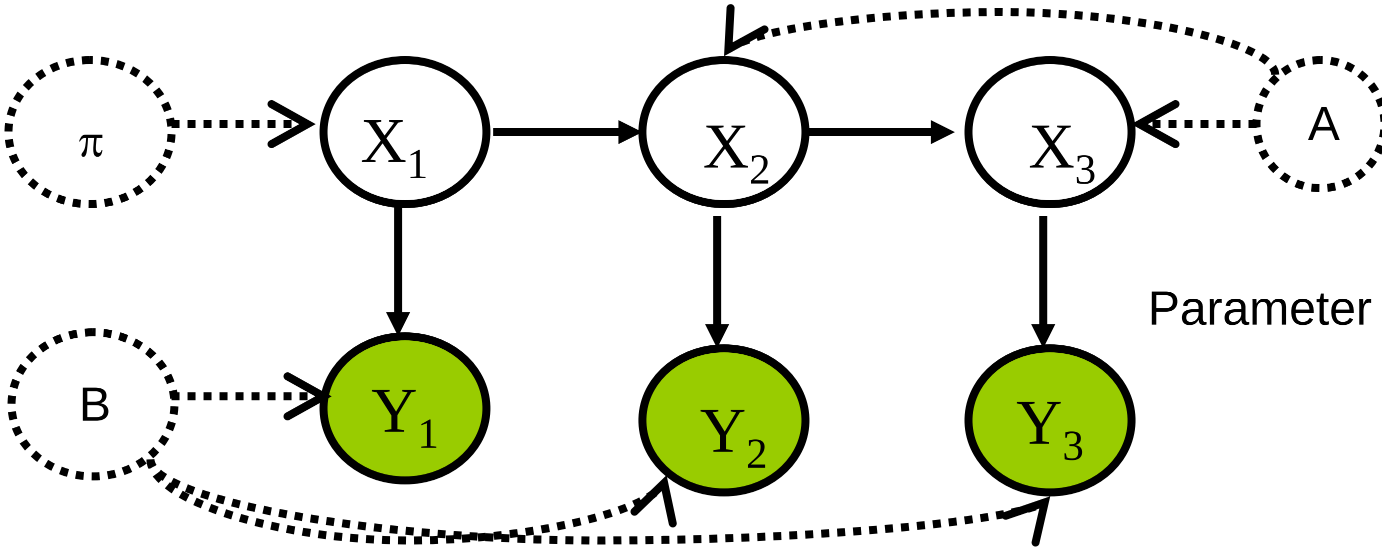

Parameter tying

$$P(X_{1:T}, Y_{1:T}) = P(X_1)P(Y_1|X_1)\prod_{t=2}^{T} P(X_t|X_{t-1})P(Y_t|X_t)$$

Transition matrix $$P(X_t = j | X_{t-1} = i) \;=\; A(i,j)$$

Observation matrix $$P(Y_t = j | X_t = i) \;=\; B(i,j)$$

Initial state distribution $$P(X_1 = i) \;=\; \pi(i)$$

Cognitive Load HMM

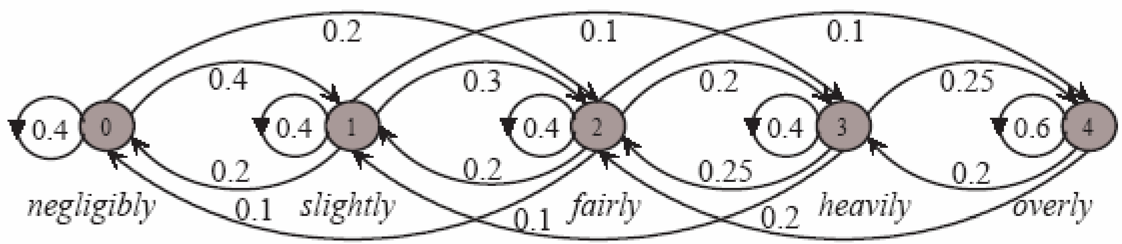

$$
B = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \\ 4 \end{array}
\begin{array}{c}
\begin{array}{cccccccccc} 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & \geq 9 \end{array} \\
\left[\begin{array}{cccccccccc}
0 & 0 & 0 & 0 & 0 & 0.02 & 0.03 & 0.05 & 0.1 & 0.8 \\
0 & 0 & 0 & 0 & 0 & 0.05 & 0.05 & 0.1 & 0.7 & 0.1 \\
0 & 0 & 0 & 0 & 0.01 & 0.02 & 0.45 & 0.4 & 0.1 & 0.02 \\
0.02 & 0.03 & 0.05 & 0.15 & 0.4 & 0.3 & 0.03 & 0.02 & 0 & 0 \\
0.1 & 0.3 & 0.3 & 0.2 & 0.1 & 0 & 0 & 0 & 0 & 0
\end{array}\right]
\end{array}
$$

Estimate hidden state (load) from observed secondary task of items remembered

Load-Sensitive Information Processing

- Depending on their load, agents will behave differently:
  - Overly: ignore incoming info
  - Heavily: randomly process half info from some of the agents and ignore all info from rest of agents
  - Fairly: process half info from any teammate
  - Slightly: process all info from some of the agents and randomly process half of the info from the rest of the agents

SMMAll

- Shared Mental Models for All: cognitive agent architecture for supporting human-centric collaborative computing
- Supported functionality includes:
  - Multi-party communication for communicating to different sections of the team
  - Shared belief map: visual display that highlights differences in information known to different team members
  - HMM models to track human processing loads

Experiment

- HAP teams must selectively share information in a timely manner to develop global situation awareness
- Humans perform 2 tasks:
  - Secondary: remember and mark belief map cells being illuminated
  - Primary: share information with right party at right time
- To decide whether to share info, humans must decide whether:
  - Info is associated with a changed belief map cell
  - Who to send the info to
  - When to send the info to

Results

- Measured mental model overlapping percentages (intersection of shared information relative to union of information)
- Demonstrated that teams that took cognitive load into account performed better at information sharing task
- HMM cognitive load model accurately tracks human cognitive load

Conclusion

- To assist human teams agents must model a wide range of factors:
  - Taskwork indicators
    - E.g. search patterns
  - Teamwork indicators
    - Communication
    - Shared mental models
  - Human’s internal state
    - Cognitive load
- Standard optimal planning techniques fall short in this domain.
- Highly active and interesting research area!