1 handout: slides

final blog entries were due

Break

- Wed May 2: HMMs, unsupervised learning, applications
- Mon May 7: special guest Scott Kiesel on robot planning
- Wed May 9, 9-noon: project presentations
- Thur May 10, 8am: paper drafts (optional for some)
- Fri May 11, 10:30: exam 3 (N133)
- Tues May 15, 3pm: papers (one hardcopy + electronic PDF)

menu?

Hidden Markov Models

MDPs:

Naive Bayes:

k-**Means:**

Markov chain:

Hidden Markov model:

A Hidden Markov Model

$$
\begin{aligned}
P(x_t = j) &= \sum_i P(x_{t-1} = i) P(x_t = j | x_{t-1} = i) \\
P(e_t = k) &= \sum_i P(x_t = i) P(e = k | x = i)
\end{aligned}
$$

More concisely:

$$
\begin{aligned}
P(x_t) &= \sum_{x_{t-1}} P(x_{t-1}) P(x_t | x_{t-1}) \\
P(e_t) &= \sum_{x_t} P(x_t) P(e | x)
\end{aligned}
$$

Viterbi Decoding

given: transition model $T(s, s')$
sensing model $S(s, o)$
observations $o_1, \ldots, o_T$
find: most probable $s_1, \ldots, s_T$

initialize $S \times T$ matrix v with 0s
$v_{0,0} \leftarrow 1$
for each time $t = 0$ to $T - 1$
    for each state s
        for each new state s'
            score $\leftarrow v_{s,t} \cdot T(s, s') \cdot S(s', o_t)$
            if score $> v_{s',t+1}$
                $v_{s',t+1} \leftarrow$ score
                best-parent(s')$\leftarrow s$
trace back from s with max $v_{s,T}$

applications

unsupervised learning: dimensionality reduction

EOLQs

- What question didn't you get to ask today?
- What's still confusing?
- What would you like to hear more about?

Please write down your most pressing question about AI and put it in the box on your way out.
Thanks!