

APPLICATIONS OF AI

APPLICATIONS

Adapted from slides kindly shared by Stuart Russell

Appreciations

◇ Libre Office

Share some of yours?

Announcements

Project P4: Ghostbusters out, due Dec 19 - but understand it before the final exam....

P4 last three questions: a bit of extra credit, if you really want to learn particle filters on your own!

Still time to participate on Piazza and improve your score there. But no “me too!” posts....

It is still an option to re-submit previous projects for 1/2 credit on incremental improvement

ROOM and DAY CHANGE: Final Exam in ECCR 200 Tue 2012-12-18 04:30 PM - 07:00 PM

Final will be closed book, bring 2 pages of notes, non-programmable calculator optional, 50% new material, see Piazza for topics, practice exams




Review for final on Wednesday

Outline


- ◇ Demo tracking
- ◇ Applications of AI
- ◇ Discussion on future impacts of AI

Credit to Dan Klein, Stuart Russell and Andrew Moore for most of today's slides


So Far: Foundational Methods



HMMs for Speech



Transitions with Bigrams



Training Counts

198015222 the first
 194623024 the same
 168504105 the following
 158562063 the world
 ...
 14112454 the door

 23135851162 the *

$$\hat{P}(\text{door}|\text{the}) = \frac{14112454}{23135851162}$$

$$= 0.0006$$

Figure from Huang et al page 618

Example: Digit Recognition

- Input: images / pixel grids
- Output: a digit 0-9
- Setup:
 - Get a large collection of example images, each labeled with a digit
 - Note: someone has to hand label all this data!
 - Want to learn to predict labels of new, future digit images
- Features: The attributes used to make the digit decision
 - Pixels: (6,8)=ON
 - Shape Patterns: NumComponents, AspectRatio, NumLoops
 - ...



0



1



2



1



??

Inverse RL: Motivation



- How do we specify a task like this?

[demo: hover / autorotate]

Autonomous Helicopter Setup



Position



On-board inertial
measurement unit (IMU)



Send out controls to
helicopter



Helicopter MDP

- **State:** $s = (x, y, z, \phi, \theta, \psi, \dot{x}, \dot{y}, \dot{z}, \dot{\psi}, \dot{\theta}, \dot{\phi})$
- **Actions (control inputs):**
 - a_{lon} : Main rotor longitudinal cyclic pitch control (affects pitch rate)
 - a_{lat} : Main rotor latitudinal cyclic pitch control (affects roll rate)
 - a_{coll} : Main rotor collective pitch (affects main rotor thrust)
 - a_{rud} : Tail rotor collective pitch (affects tail rotor thrust)
- **Transitions (dynamics):**
 - $s_{t+1} = f(s_t, a_t) + w_t$
[f encodes helicopter dynamics]
[w is a probabilistic noise model]
- **Can we solve the MDP yet?**



Problem: What's the Reward?

- Rewards for hovering:

[demo: hover / tic-toc]

$$R(s) = -(\alpha_x(x - x^*)^2 + \alpha_y(y - y^*)^2 + \alpha_z(z - z^*)^2 + \alpha_{\dot{x}}(\dot{x} - \dot{x}^*)^2 + \alpha_{\dot{y}}(\dot{y} - \dot{y}^*)^2 + \alpha_{\dot{z}}(\dot{z} - \dot{z}^*)^2)$$

- Rewards for “Tic-Toc”?

- Problem: what's the target trajectory?
- Just write it down by hand?



[demo: bad]

Apprenticeship Learning

- Goal: learn reward function from expert demonstration
- Assume $R(s) = w \cdot f(s)$
- Get expert demonstrations $\mathbf{s} = (s_0, s_1, \dots, s_n)$
- Guess initial policy π_0
- Repeat:
 - Find w which make the expert better than $\{\pi_0, \pi_1, \dots, \pi_{i-1}\}$
 $w_i \leftarrow \text{distinguish}(\pi^*, \{\pi_0, \pi_1, \dots, \pi_{i-1}\})$
 - Solve MDP for new weights w :
 $\pi_i \leftarrow \text{solve}(MDP(w_i))$

Pacman Apprenticeship!

- Demonstrations are expert games




- Features defined over states s
- Score of a state given by:

$$w \cdot f(s)$$

- Learning goal: find weights which explain expert actions

Helicopter Apprenticeship?



Probabilistic Alignment

Intended trajectory


$$z_{t+1} = f(z_t) + \omega_t$$

Expert

demonstrations


$$y_j = z_{\tau_j} + \nu_j$$

Time indices



- Intended trajectory satisfies dynamics.
- Expert trajectory is a noisy observation of one of the hidden states.
 - But we don't know exactly which one.

Alignment of Samples



- Result: inferred sequence is much cleaner!

Final Behavior



What is NLP?



- Fundamental goal: analyze and process human language, broadly, robustly, accurately...
- End systems that we want to build:
 - Ambitious: speech recognition, machine translation, information extraction, dialog interfaces, question answering...
 - Modest: spelling correction, text categorization...



Problem: Ambiguities

- **Headlines:**

- Enraged Cow Injures Farmer With Ax
- Hospitals Are Sued by 7 Foot Doctors
- Ban on Nude Dancing on Governor's Desk
- Iraqi Head Seeks Arms
- Local HS Dropouts Cut in Half
- Juvenile Court to Try Shooting Defendant
- Stolen Painting Found by Tree
- Kids Make Nutritious Snacks

- **Why are these funny?**



Parsing as Search



Hershey bars protest


Grammar: PCFGs

- Natural language grammars are very ambiguous!
- PCFGs are a formal probabilistic model of trees
 - Each “rule” has a conditional probability (like an HMM)
 - Tree’s probability is the product of all rules used
- Parsing: Given a sentence, find the best tree – search!



ROOT → S	375/420
S → NP VP .	320/392
NP → PRP	127/539
VP → VBD ADJP	32/401
.....	

Syntactic Analysis



Hurricane Emily howled toward Mexico 's Caribbean coast on Sunday packing 135 mph winds and torrential rain and causing panic in Cancun, where frightened tourists squeezed into musty shelters .

Machine Translation

"Il est impossible aux journalistes de rentrer dans les régions tibétaines"

Bruno Philip, correspondant du "Monde" en Chine, estime que les journalistes de l'AFP qui ont été expulsés de la province tibétaine du Qinghai "n'étaient pas dans l'illégalité".

Les faits Le dalaï lama dénonce l'"enfer" imposé au Tibet depuis sa fuite, en 1959

Vidéo Anniversaire de la rébellion tibétaine : la Chine en garde



"It is impossible for journalists to enter Tibetan areas"

Philip Bruno, correspondent for "World" in China, said that journalists of the AFP who have been deported from the Tibetan province of Qinghai "were not illegal."

Facts The Dalai Lama denounces the "hell" imposed since he fled Tibet in 1959

Video Anniversary of the Tibetan rebellion: China on guard

Real-time Translation



- Translate text from one language to another
- Recombines fragments of example translations
- Challenges:
 - What fragments? [learning to translate]
 - How to make efficient? [fast translation search]



The Problem with Dictionary Look-ups

顶部	/ top /roof/
顶端	/summit/peak/ top /apex/
顶头	/coming directly towards one/ top /end/
盖	/lid/ top /cover/canopy/build/Gai/
盖帽	/surpass/ top /
极	/extremely/pole/utmost/ top /collect/receive/
尖峰	/peak/ top /
面	/fade/side/surface/aspect/ top /face/flour/
摘心	/ top /topping/

Example from Douglas Hofstadter

A Brief and Biased History



Warren Weaver

When I look at an article in Russian, I say: "This is really written in English, but it has been coded in some strange symbols. I will now proceed to decode."



John Pierce

"Machine Translation" presumably means going by algorithm from machine-readable source text to useful target text... In this context, there has been no machine translation...

Berkeley's first MT grant

Statistical MT thrives

MT is the "first" non-numeral compute task

ALPAC report deems MT bad

Statistical data-driven approach introduced

