


CS 188: Artificial Intelligence Spring 2010

Lecture 27: Conclusion
4/28/2010

Pieter Abbeel – UC Berkeley

Announcements

- Project 5 due tonight.



- Office hours
 - next week: only Woody and Alex.
 - Next next week: back to normal office hours.
- Contest!!
 - Tournaments every night.
 - Final tournament: We will use submissions received by Thursday May 6, 11pm.**

Today

- AI applications:
 - Robotics, language, vision



- Where to go next

3

Quadruped



- Model as MDP --- but what should the reward function be?
- Reward function trades off 25 features.
- $R(x) = w^T \phi(x)$

[Kolter, Abbeel & Ng, 2008]

Find the Reward Function for Foot Placements

- $R(x) = w^T f(x)$
- Find the reward function that makes the demonstrations better than all other paths by some margin

$$\min_{w, \xi \geq 0} \|w\|_2^2 + \sum_i \xi^{(i)}$$

$$\text{s.t.} \quad \forall i, \forall x, \sum_t w \cdot f(x_t^{(i)}) \geq \sum_t w \cdot f(x_t) + 1 - \xi^{(i)}$$

- Compare to support vector machine (SVM):


$$\min_{w, \xi \geq 0} \|w\|_2^2 + \sum_i \xi^{(i)}$$

$$\text{s.t.} \quad \forall i, y \cdot w_y \cdot f(x_i) \geq w_y \cdot f(x_i) + 1 - \xi^{(i)}$$


5

High-Level Control

- Demonstrate path across the "training terrain"



- Run apprenticeship learning to find a set of weights w
- Receive "testing terrain" (a height map)



- Find a policy for crossing the testing terrain.

Without learning



With learned reward function

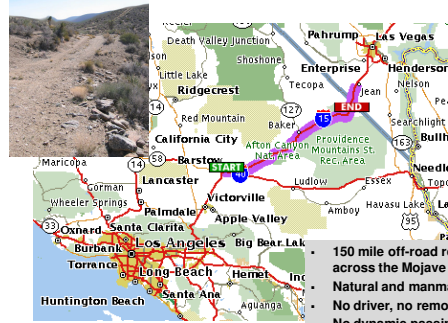


Autonomous Vehicles




Autonomous vehicle slides adapted from Sebastian Thrun

Grand Challenge: Barstow, CA, to Primm, NV




- 150 mile off-road robot race across the Mojave desert
- Natural and manmade hazards
- No driver, no remote control
- No dynamic passing

Inside an Autonomous Car



Sensors: Camera



Vision for a Car




Self-Supervised Vision




- Demo: lidar-and-vision 1.mp4

16

Readings: No Obstacles




Readings: Obstacles




Obstacle Detection

Trigger if $|z^i - z^j| > 1.5\text{cm}$ for nearby z^i, z^j




Raw Measurements: 12.6% false positives


Probabilistic Error Model



HMMs for Detection



Raw Measurements: 12.6% false positives HMM Inference: 0.02% false positives



Demo: Urban challenge intersection

What is NLP?



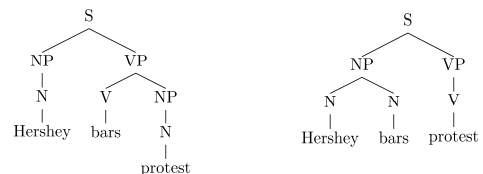
- Fundamental goal: analyze and process human language, broadly, robustly, accurately...
- End systems that we want to build:
 - Ambitious: speech recognition, machine translation, information extraction, dialog interfaces, question answering...
 - Modest: spelling correction, text categorization...

27

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

29


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!



30

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.

32

Machine Translation

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

Source: Philip, un correspondant de Reuters en Chine, annonce qu'une mission de journalistes de l'AFP est empêchée de pénétrer dans les régions tibétaines de Chine.

"It is impossible for journalists to enter Tibetan areas"

Philip Thomas, correspondent for Reuters in China, said that journalists of the AFP who have been expelled from the Tibetan province of Lhasa "may not return."

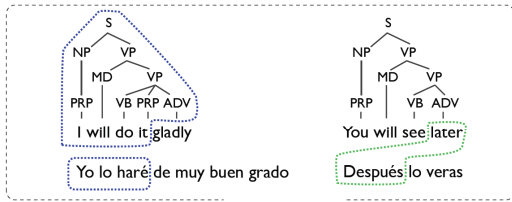
- Translate text from one language to another
- Recombines fragments of example translations
- Challenges:
 - What fragments? [learning to translate]
 - How to make efficient? [fast translation search]

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

Machine Translation



Machine translation system:




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...



Data-Driven Machine Translation

Target language corpus:

I will get to it soon See you later He will do it

Sentence-aligned parallel corpus:

Yo lo haré mañana Hasta pronto Hasta pronto
 I will do it tomorrow See you soon See you around

Machine translation system:

Yo lo haré pronto (NOVEL SENTENCE) → Model of translation → I will do it soon

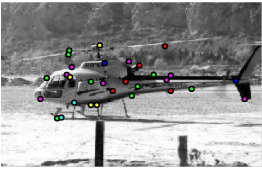

Learning to Translate

CLASSIC SOUPS

	Sm.	Lg.
水 推 雞 湯 57. House Chicken Soup (Chicken, Celery, Potato, Onion, Carrot)	1.50	2.75
雞 飯 湯 58. Chicken Rice Soup	1.85	3.25
雞 麵 湯 59. Chicken Noodle Soup	1.85	3.25
廣 東 湯 60. Cantonese (Wonton) Soup	1.50	2.75
蕃 茄 湯 61. Tomato Clear Egg Drop Soup	1.65	2.95
餃 子 湯 62. Regular (Wonton) Soup	1.10	2.10
酸 辣 湯 63. Hot & Sour Soup	1.10	2.10
香 花 湯 64. Egg Drop Soup	1.10	2.10
雲 吞 湯 65. Egg Drop (Wonton) Mix	1.10	2.10
豆 腐 菜 湯 66. Tofu Vegetable Soup	NA	3.50
雞 王 米 湯 67. Chicken Corn Cream Soup	NA	3.50
蟹 肉 玉 米 湯 68. Crab Meat Corn Cream Soup	NA	3.50
海 鮮 湯 69. Seafood Soup	NA	3.50

Example from Adam Lopez

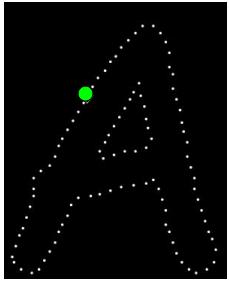
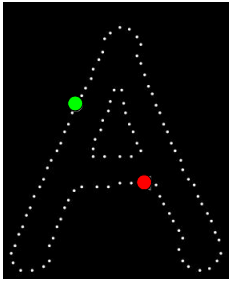
Computer Vision: Object Recognition

Template Query

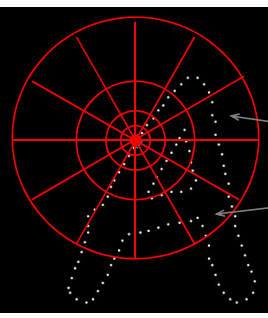
44

Comparing Local Regions

45

Shape Context



Count the number of points inside each bin, e.g.:

Count = 4

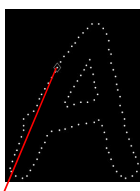
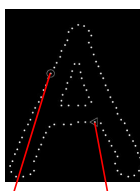
⋮

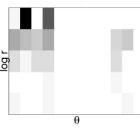
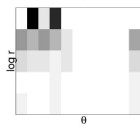
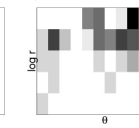
Count = 10

◆ Compact representation of distribution of points relative to each point

46

Shape Context

47

Incorporation into nearest neighbors

- Start from two images --- extract contour points
- Compute pairwise cost matrix
 - Cost for difference in histogram
 - Cost difference in tangent direction
- Find lowest cost 1-to-1 matching between two images
- Find the geometric transformation as close as possible to this 1-to-1 matching
- Distance between two images consists of
 - Cost of each of the pairs in the 1-to-1 matching
 - Closeness of geometric transformation and the 1-to-1 matching

50

Empirical Evaluation

- Tested their approach on the MNIST dataset.
- As of Dec 2009, more than 50 algorithms have been tested on the database.
- Training set: 60,000 examples.
- Test set: 10,000 examples.
- Error rate: 0.63% using 20,000 training examples and 3-NN.
- At the time of publication, this error rate was the lowest. Currently, the lowest error rate is 0.39%.

51

Where to go next?

- Congratulations, you've seen the basics of modern AI
 - ... and done some amazing work putting it to use!
- How to continue:
 - Robotics / vision / IR / language: cs189
 - Machine learning: cs281a / cs281b
 - Cognitive modeling: cog sci 131
 - Vision: cs280
 - Robotics: cs287
 - NLP: cs288
 - Optimization: ee127a and ee227a
 - ... and more; ask if you're interested

52

That's It!

- Happy studying, good luck on the exam and contest, and have a great summer!

53