CS 188: Artificial Intelligence

Conclusion

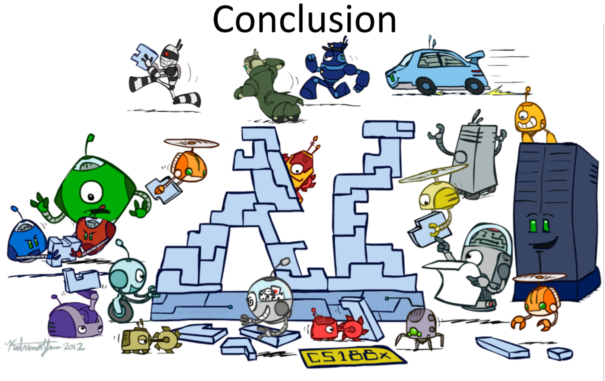

Instructors: Dan Klein and Pieter Abbeel --- University of California, Berkeley

Contest Results

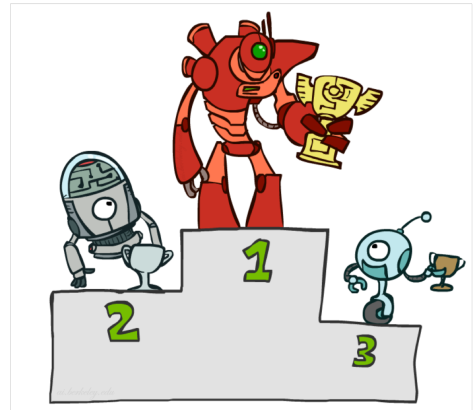

P1 Mini-Contest Results!

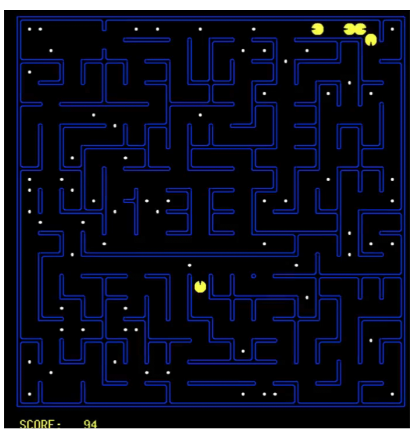

P1 Mini-Contest Results

- 1st place:
  - Yunsheng — Yunsheng Ma, Ryan Xie
- 2nd place
  - JasonL — Jingyuan Li
- 3rd place
  - Winnie-the-Pooh — Philip Zhao, Winnie Gao

P2 Mini-Contest Results!

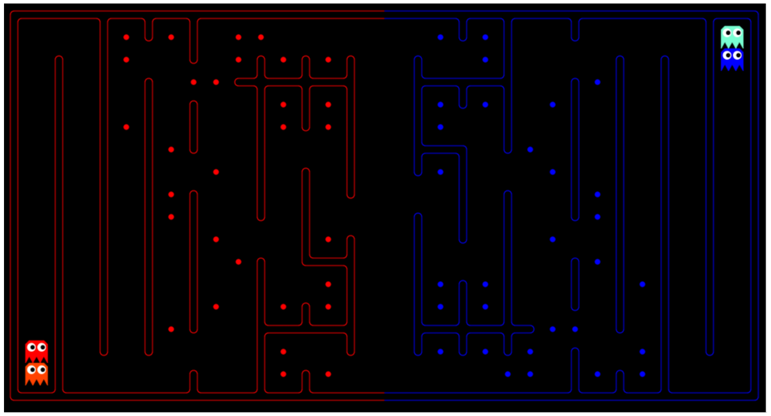

P2 Mini-Contest Results

- 1st place:
  - @_@ Philip Zhao, Winnie Gao
- 2nd place
  - YZWY Yuechen Wu, Yuzishu
- 3rd place
  - DON'T FORGET: REGISTER TO VOTE! Sean Liu, Ham Huang

Final Contest!

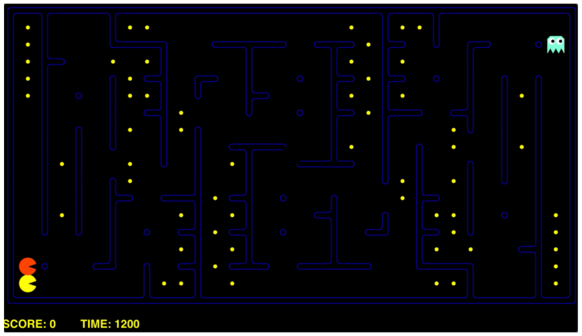

A cooperative version of PacMan where you write a bot to coordinate with another bot to gather food and defeat ghosts.

Final Contest Statistics

- 32 teams, thousands of matches!
- Great work by everyone!
- Creative Names:
  - pacmantaughtmelife
  - Stupid Pacman is not Ready
  - broken bot
  - Basic bot
  - Ôº∅Ôº¢Ôº≥Ôº•Ôº≤Ôº∂Ôº•Ôº≤
  - He never listens
  - SPAM
  - Pacman is READY!!
  - debug_fixed?
- Final results: now

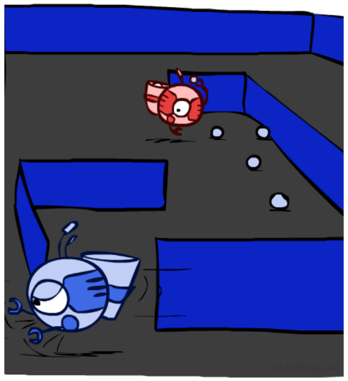

Top-10

4 [757]	Watney The Fearful	Alexander Khazatsky
5 [761]	debug_fixed?	mssheldonmao
6 [802]	nine (9) v3	Wilson Wu
7 [812]	DieGhostDie	Shi Mao, Zhibo Fan
8 [813]	openai five candidate	Martin Li
9 [817]	First attempt v4.6	Fredrik Roemming
10 [848]	Mark-??	Winnie Gao, Philip Zhao

3rd Place – WhenMonaSmiles – Victor Cheng

- The bot is based on ReflexCaptureAgent using an using an feature based evaluation function. The features are teammateDistance, distanceToFood, and ghostDistance.
- Basically, the agent aims for the furthest food when it is a certain distance close to the teammate, otherwise it aims for the closest food.
- The agent tries to get away from the ghost when it is close to the ghost.
- The agent values food more than the danger of ghost, as getting the attention of the ghost would potentially help the teammate.
- The thresholds of distances to the teammate and ghost, need to be tuned, like the weights of the feature. Optimally they would be tuned by RL or other learning methods, but they are tuned manually this time.

3rd Place – WhenMonaSmiles – Victor Cheng

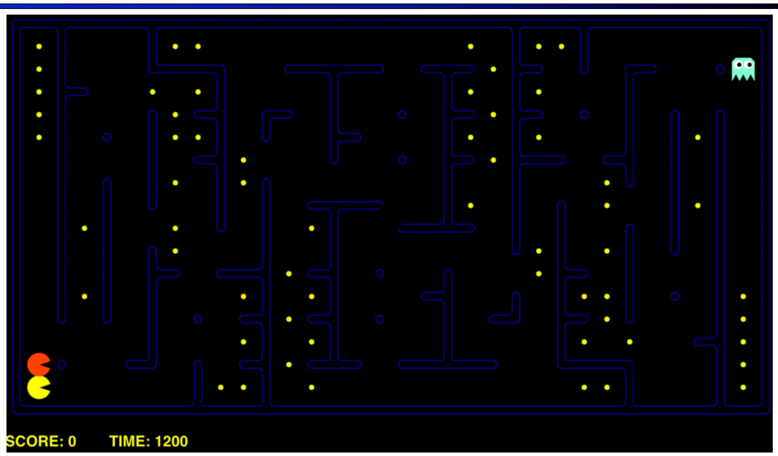

2nd Place – Yihe Huang

- Strategy is based on Approximate Q-Learning.
- Features include 3 distances and 2 scores.
- Distances include the maze distance between ghost, teammate, nearest food and my bot.
- Scores include the successor score and a score for exploring and exploiting (to avoid deadlock).

2nd Place – Yihe Huang

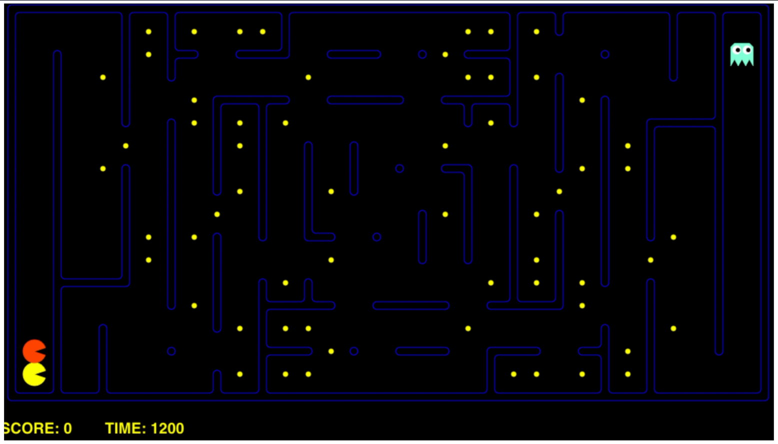

1st Place – Rudy Zhang & Feng Xu

- Strategy is based on a map named reward density. It is calculated in the following steps:
  - 1. calculate food density like minesweepers
  - 2. lower the reward of the area if the teammate might approach the area using particle filtering to update teammate position beliefs
  - 3. adjust the reward of a position according to the Pacman's distance to it
  - 4. lower the reward of a position if a ghost is near it
- Using the computed reward density map, we have the following strategy
  - If the ghost is not nearby:
    - go to the position with maximum reward density then collect the food local optimally (both using star search)
  - else if the ghost is close:
    - use minimax strategy to avoid the ghost, award the Pacman for approaching the max reward density position
- Some special calculation:
  - cached actions from the start position to the first position with len(legal actions)>1

1st Place – Rudy Zhang & Feng Xu

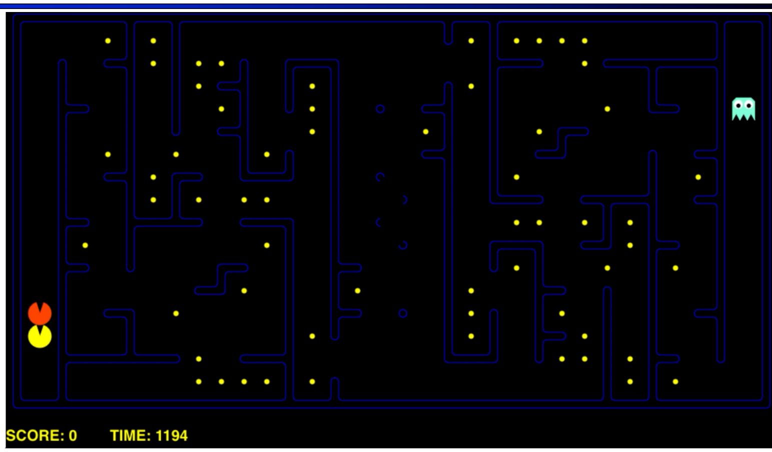

Top-10

1 [711]	Pacman No.70 is READY	Rudy Zhang, Feng Xu
2 [720]	MasterYi	Yihe Huang
3 [748]	WhenMonaSmiles	Victor Cheng
4 [757]	Watney The Fearful	Alexander Khazatsky
5 [761]	debug_fixed?	mssheldonmao
6 [802]	nine (9) v3	Wilson Wu
7 [812]	DieGhostDie	Shi Mao, Zhibo Fan
8 [813]	openai five candidate	Martin Li
9 [817]	First attempt v4.6	Fredrik Roemming
10 [848]	Mark-??	Winnie Gao, Philip Zhao

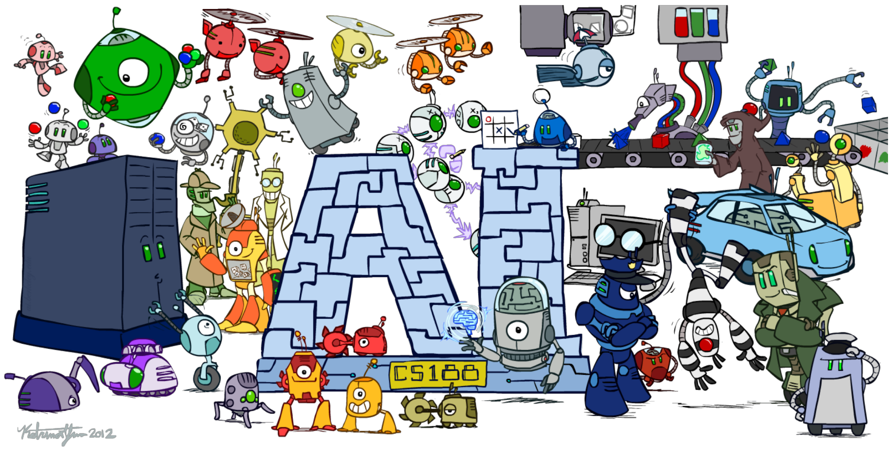

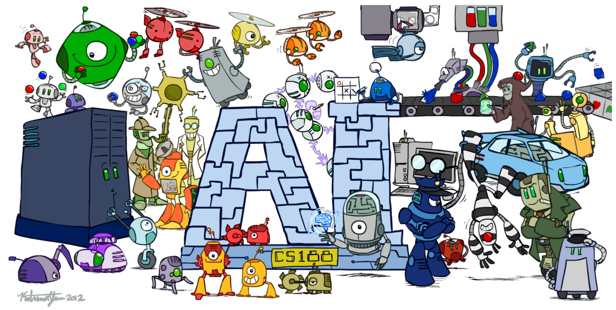

Ketrina Yim
CS188 Artist

Language Technologies

Goal: Deep Understanding

- Requires context, linguistic structure, meanings…

Reality: Shallow Matching

- Requires robustness and scale
- Amazing successes, but fundamental limitations

Speech Recognition in an ~~Hour~~ 15 Min

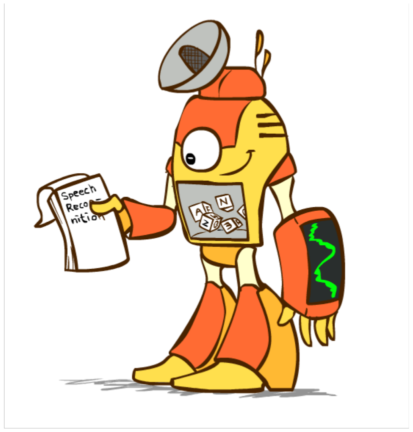

Why is Speech Recognition Hard?

Digitizing Speech

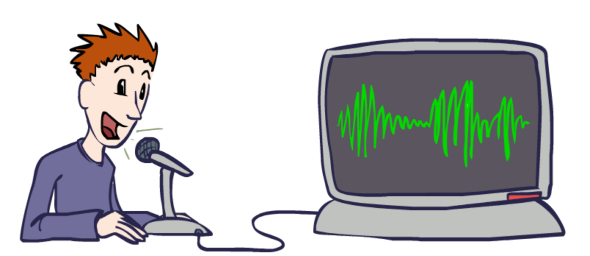

Speech Input

Speech input is an acoustic waveform

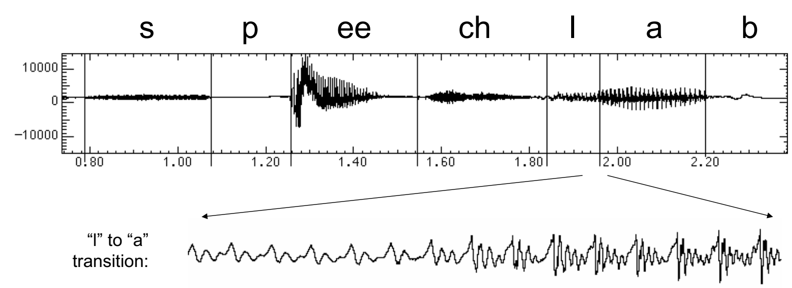

Figure: Simon Arnfield, http://www.psyc.leeds.ac.uk/research/cogn/speech/tutorial/

Part of [ae] from “lab”

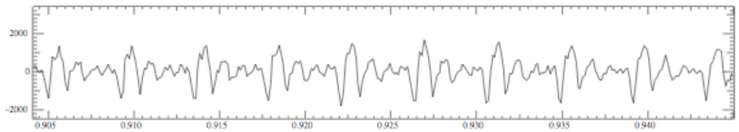

- Complex wave repeating nine times
  - Plus smaller wave that repeats 4x for every large cycle
  - Large wave: freq of 250 Hz (9 times in .036 seconds)
  - Small wave roughly 4 times this, or roughly 1000 Hz

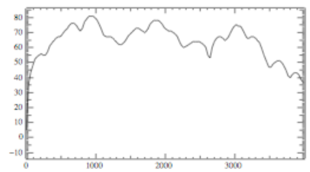

Spectral Analysis

- Frequency gives pitch; amplitude gives volume
  - Sampling at ~8 kHz (phone), ~16 kHz (mic) (kHz=1000 cycles/sec)

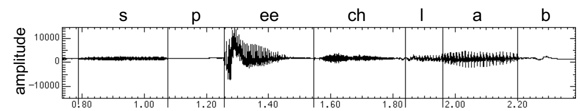

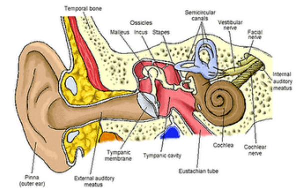

- Fourier transform of wave displayed as a spectrogram
  - Darkness indicates energy at each frequency

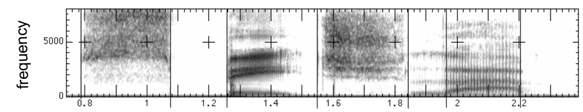

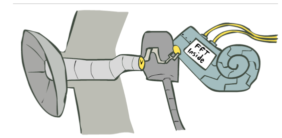

Human ear figure: depion.blogspot.com

Why These Spectral Peaks?

- Articulator process:
  - Vocal cord vibrations create harmonics
  - The mouth is an amplifier
  - Depending on shape of mouth, some harmonics are amplified more than others

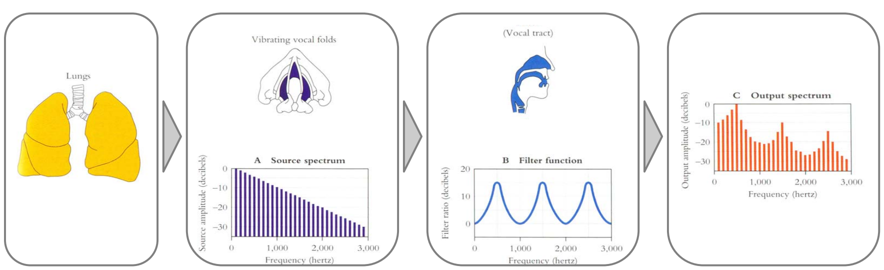

Resonances of the Vocal Tract

- The human vocal tract as an open tube

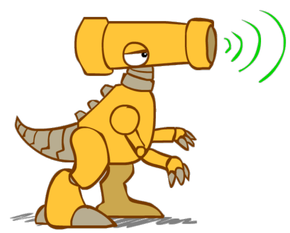

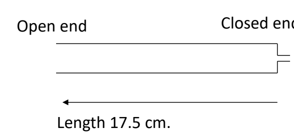

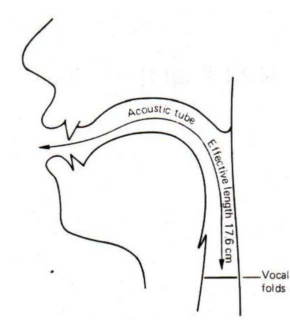

- Air in a tube of a given length will tend to vibrate at resonance frequency of tube
- Constraint: Pressure differential should be maximal at (closed) glottal end and minimal at (open) lip end

Figure: W. Barry Speech Science slides

Spectrum Shapes

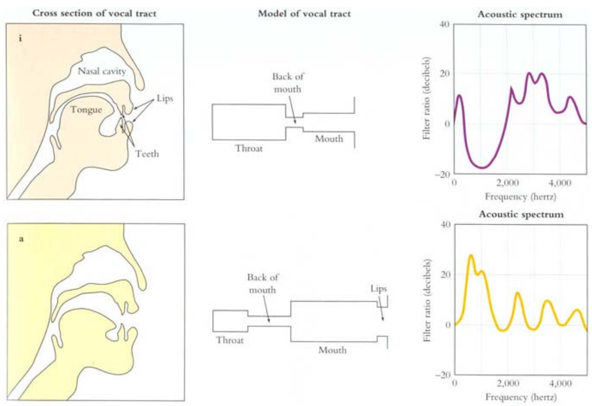

[demo]

Figure: Mark Liberman

Vowel [i] Sung at Successively Higher Pitches

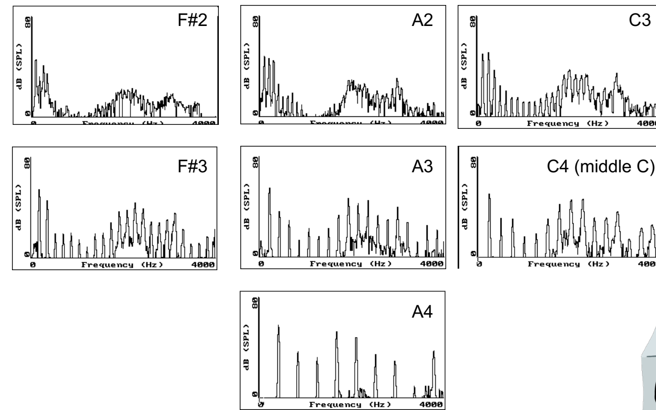

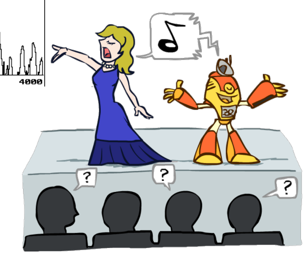

Graphs: Ratree Wayland

Speech Recognition as an HMM

- Evidence: Sequences of acoustic vectors (~39 real numbers per slice)

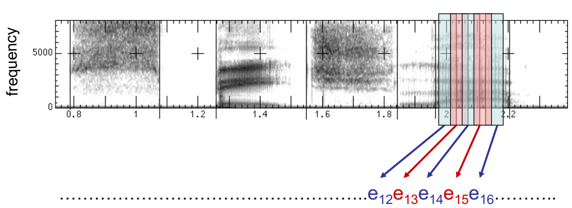

- Hidden states: Which words were spoken? Almost!

Speech Recognition State Space

- HMM Specification
  - P(E|X) models which acoustic vectors match each phoneme (each kind of sound)
  - P(X|X') encodes how sounds can be strung together
- State Space
  - We will have one state for each sound in each word ("pronunciation cursor")
  - Mostly, states advance sound by sound along a word
  - We build a little state graph for each word and chain them together to form the state space X

States in a Word

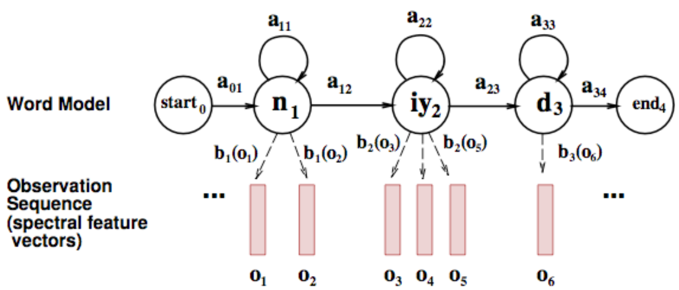

Transitions with a Bigram Language Model

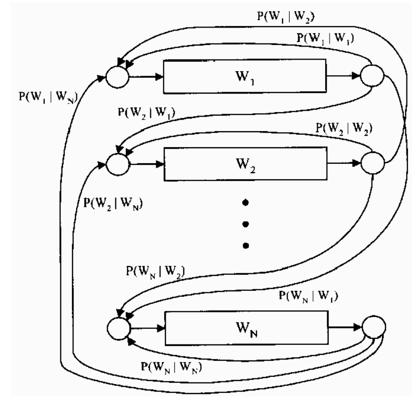

Training Counts

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: Huang et al, p. 618

Increasing N-Gram Order

More history captures more correlations

Bigram Model

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

23135851162 the *	

Trigram Model

197302	close the window
191125	close the door
152500	close the gap
116451	close the thread
87298	close the deal

3785230 close the *	

P(door | the) = 0.0006

P(door | close the) = 0.05

Decoding

- Finding the words given the acoustics is an HMM inference problem
- Which state sequence $x_{1:T}$ is most likely given the evidence $e_{1:T}$?

$$x^*_{1:T} = \underset{x_{1:T}}{\arg\max}\, P(x_{1:T}|e_{1:T}) \;= \underset{x_{1:T}}{\arg\max}\, P(x_{1:T}, e_{1:T})$$

- From the sequence x, we can simply read off the words

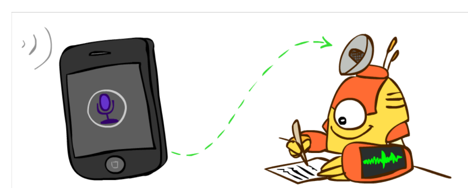

Neural Nets for Speech

- Major advances in ASR over the last ~5 years due to neural nets
  - Acoustic models P(frequencies | phones) now parameterized with NNs
  - Language models P(word | word history) now parameterized with NNs

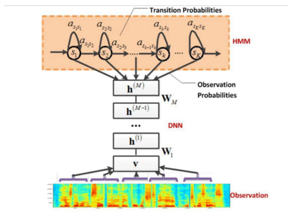

Dahl, George E., et al. "Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition." *Audio, Speech, and Language Processing, IEEE Transactions on* 20.1 (2012): 30-42.

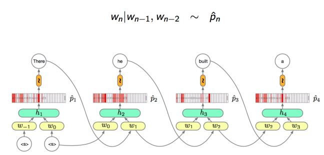

Figure: 52coding.com

Pac-Man Beyond the Game!

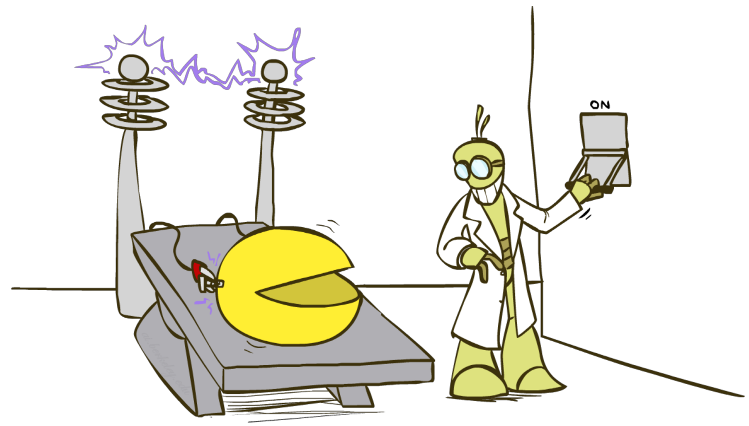

Pacman: Beyond Simulation?

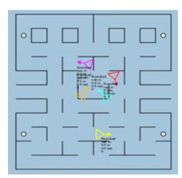

Students at Colorado University: http://pacman.elstonj.com

[VIDEO: Roomba Pacman.mp4]

Pacman: Beyond Simulation!

Bugman?

- AI = Animal Intelligence?
  - Wim van Eck at Leiden University
  - Pacman controlled by a human
  - Ghosts controlled by crickets
  - Vibrations drive crickets toward or away from Pacman's location

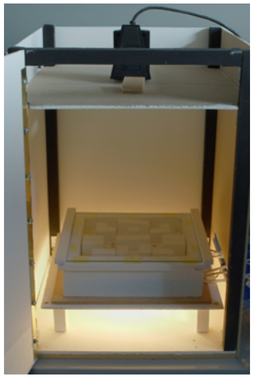

http://pong.hku.nl/~wim/bugman.htm

Bugman

[VIDEO: bugman_movie_1.mov]

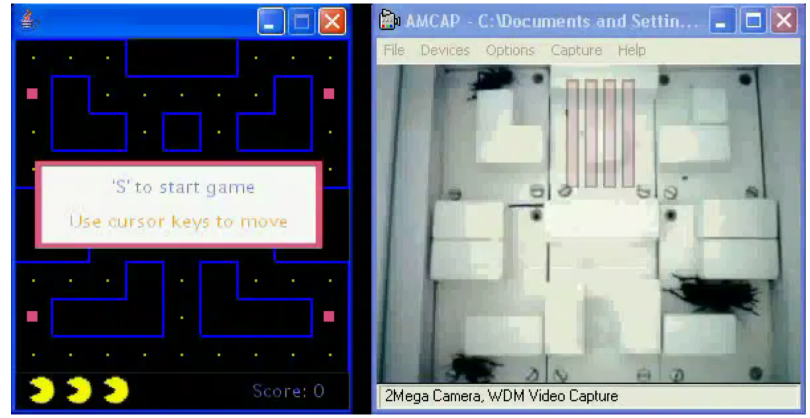

Crawler

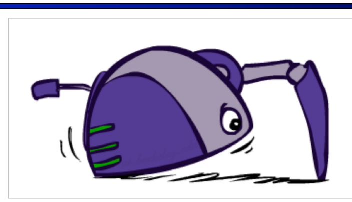

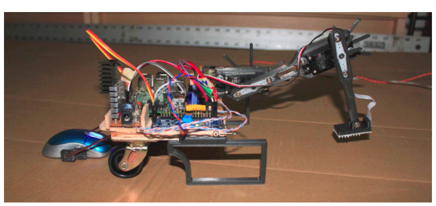

Q-learning with Robot Crawler

[Jan Balewski]

Where to Go Next?

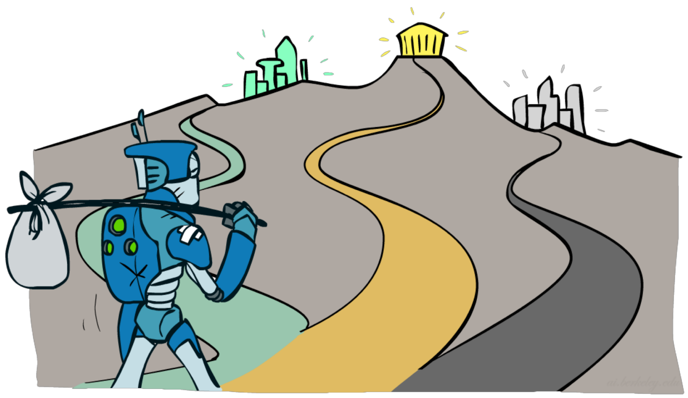

Where to go next?

- Congratulations, you've seen the basics of modern AI
  - … and done some amazing work putting it to use!
- How to continue:
  - Machine learning: cs189, cs182, stat154
  - Intro to Data Science: Data 100
  - Probability: ee126, stat134
  - Optimization: ee127
  - Cognitive modeling: cog sci 131
  - Machine learning theory: cs281a/b
  - Vision: cs280
  - Robotics: cs287
  - Algorithmic Human Robot Interaction: cs294-115
  - Reinforcement Learning: cs285
  - NLP: cs288
  - … and more; ask if you're interested

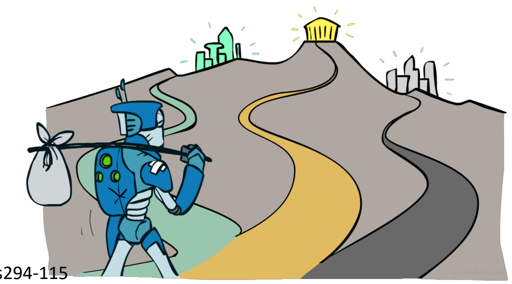

How about AI Research?

https://bair.berkeley.edu

That's It!

- Help us out with some course evaluations
- Have a great summer, and always maximize your expected utilities!