CSE 473

Final Lecture: A Smörgåsbord of AI Applications

AI in the today's news!

The New York Times

Tuesday, December 5, 2006 Last Update: 10:11 PM ET

Once Again, Machine Beats Human Champion at Chess

By DYLAN LOEB McCLAIN
Published: December 5, 2006

E-MAIL
PRINT
SAVE

ARTICLE TOOLS SPONSORED BY

In the continuing quest to see if humans can outpace their electronic creations, the humans have lost another, perhaps decisive, round.

Henning Kaiser/AFP – Getty Images

A six-game chess match between Vladimir Kramnik of Russia, the world champion, and Deep Fritz, a souped-up version of commercially available chess software made by Chessbase, ended today in victory for the computer, which won the final game and clinched the match, 4 games to 2.

2

Deep Fritz: Details

- Based on commercially available Deep Fritz 10 software from Chessbase (selling at $137.47)
- No specialized chess hardware as in IBM's Deep Blue
- Multi-threaded
- Different search techniques

    Null-move heuristic added to α-β pruning: Skip a turn and do a shallow search

Player	Rating	1	2	3	4	5	6	Tot.
Vladimir Kramnik	2760	½	0	½	½	½	0	2.0
Deep Fritz 10	-	½	1	½	½	½	1	4.0

3

Recall: Neural Networks

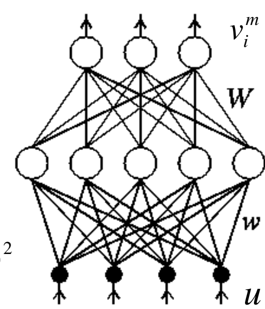

$$v_i^m = g(\sum_j W_{ji} x_j)$$

$$x_j^m = g(\sum_k w_{kj} u_k^m)$$

$$E(\mathbf{W}, \mathbf{w}) = \frac{1}{2} \sum_i (d_i - v_i)^2$$

$$u_k^m$$

Backprop rule for <u>input-hidden weights **w**</u>:

$$w_{kj} \rightarrow w_{kj} - \varepsilon \frac{dE}{dw_{kj}}$$

$$\frac{dE}{dw_{kj}} = \left[-\sum_{m,i} (d_i^m - v_i^m) g'(\sum_j W_{ji} x_j^m) W_{ji} \right] \cdot \left[g'(\sum_k w_{kj} u_k^m) u_k^m \right]$$

4

Application: Pattern Recognition

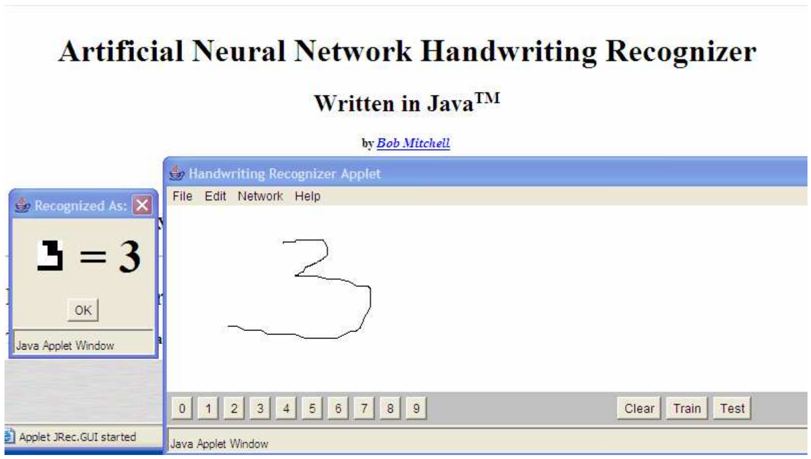

http://www-cse.uta.edu/%7Ecook/ai1/lectures/applets/hnn/JRec.html

5

Recall: Recursive Bayesian Updating

$$P(x \mid z_1,\ldots,z_n) = \frac{P(z_n \mid x, z_1,\ldots,z_{n-1})\, P(x \mid z_1,\ldots,z_{n-1})}{P(z_n \mid z_1,\ldots,z_{n-1})}$$

Markov assumption: z_n is independent of $z_1,\ldots,z_{n-1}$ if we know x.

$$P(x \mid z_1,\ldots,z_n) = \frac{P(z_n \mid x, z_1,\ldots,z_{n-1})\, P(x \mid z_1,\ldots,z_{n-1})}{P(z_n \mid z_1,\ldots,z_{n-1})}$$

$$= \frac{P(z_n \mid x)\, P(x \mid z_1,\ldots,z_{n-1})}{P(z_n \mid z_1,\ldots,z_{n-1})}$$

$$= \alpha\, P(z_n \mid x)\, P(x \mid z_1,\ldots,z_{n-1})$$

Recursive!

6

Application: Robot Localization and Mapping of Allen Center

(Work of Prof. Dieter Fox and students)

7

Recall: Bayesian Networks

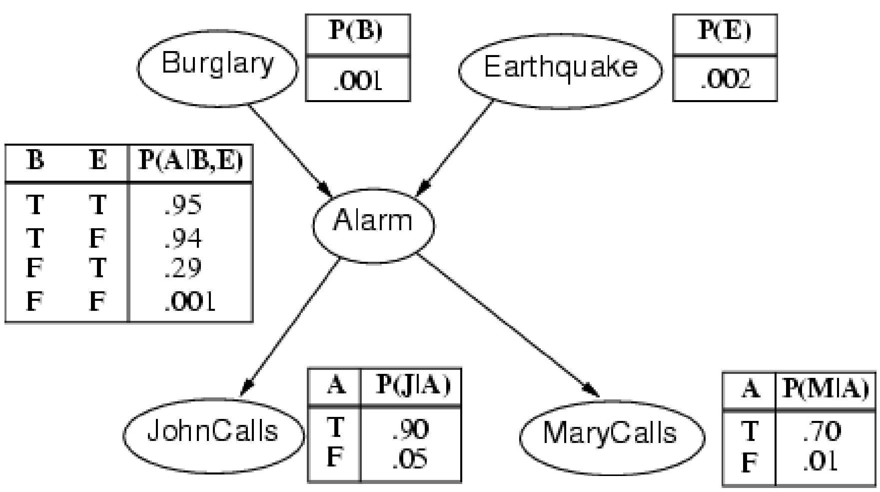

P(B)
.001

P(E)
.002

B	E	P(A|B,E)
T	T	.95
T	F	.94
F	T	.29
F	F	.001

A	P(J|A)
T	.90
F	.05

A	P(M|A)
T	.70
F	.01

8

Application: Tracking a Person using GPS

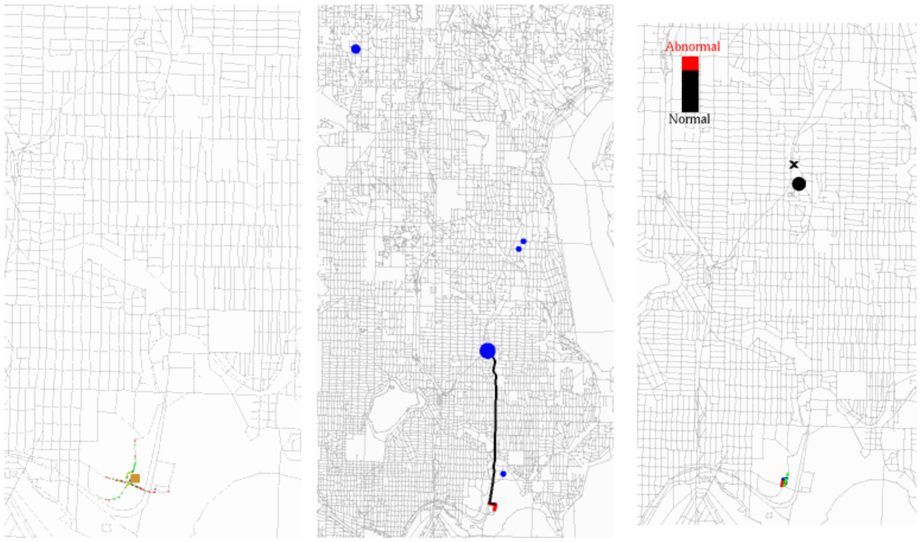

foot=blue, bus=green, car=red

(Work of Prof. Fox, Prof. Kautz, and students)

9

Application: Robot Learning by Imitation

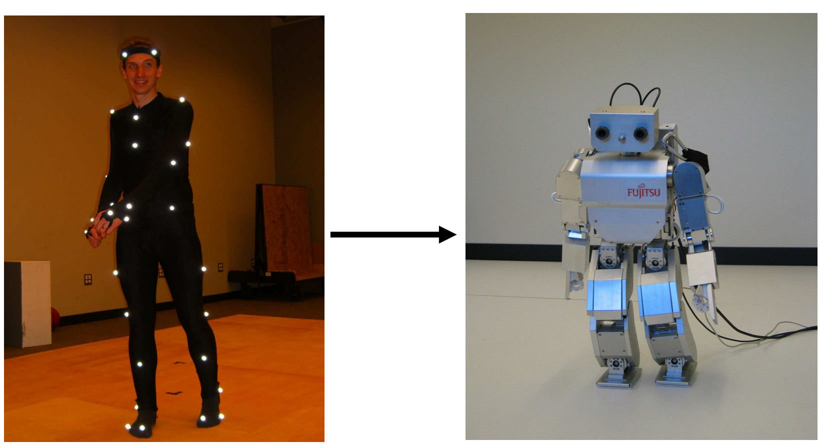

(grad student David Grimes)

10

Imitating from Motion Capture Data

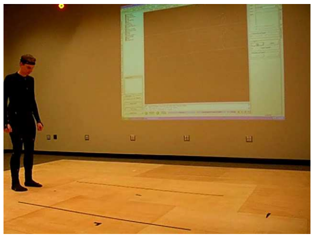

Motion Capture

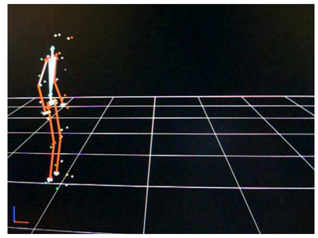

Data from Motion Capture

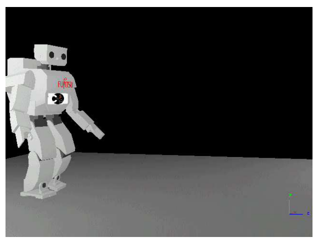

Attempted Imitation

11

Bayesian Network for Stable Imitation and Learning

Idea: Use Bayesian network to capture consequences of actions
(current body state, action) $\rightarrow$ Next body state

State s = [joint angles, gyro values, foot pressure values]

Action a = [position commands to motors for each joint]

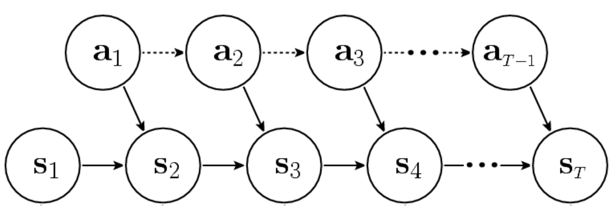

Infer actions a_t given evidence $s_1,\ldots,s_T$ from teacher subject to stability constraints on gyro readings

12

Learning to Imitate a Human Action

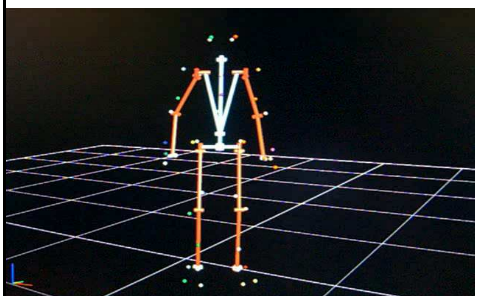

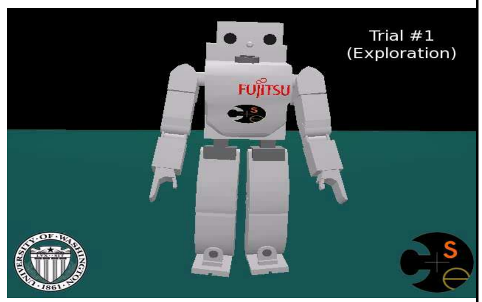

13

Result after Learning

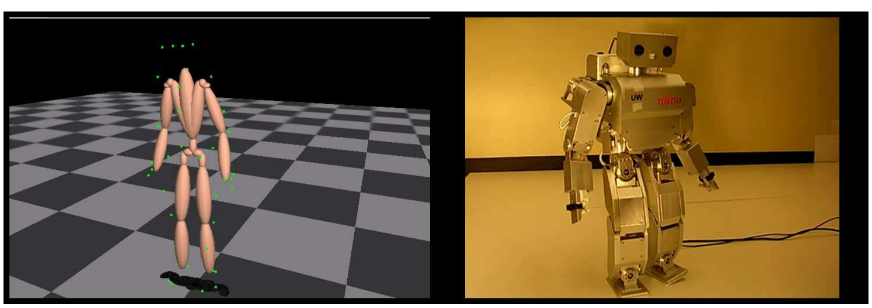

Human Action *Imitation*

14

Recall: Classification Techniques

- Decision Trees
- Nearest Neighbors
- Neural Networks
- Etc.

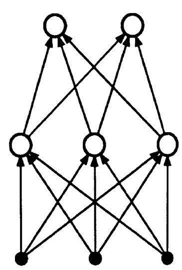

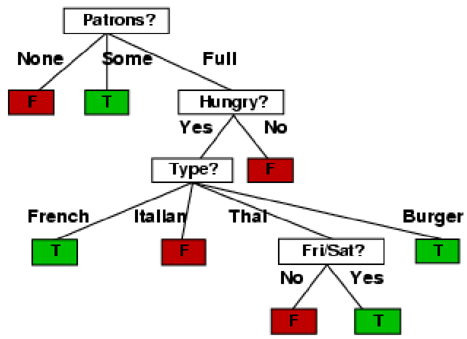

15

Application: Brain-Computer Interfaces

- Classifying brain signals recorded at the scalp
- Detect which object a person wants from a set of objects

(grad student Pradeep Shenoy)

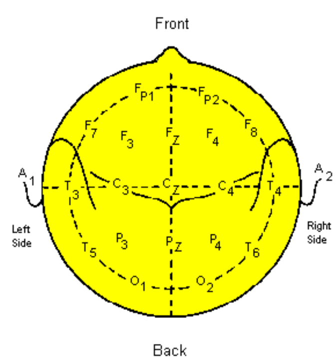

16

Deciphering choice from brain signals

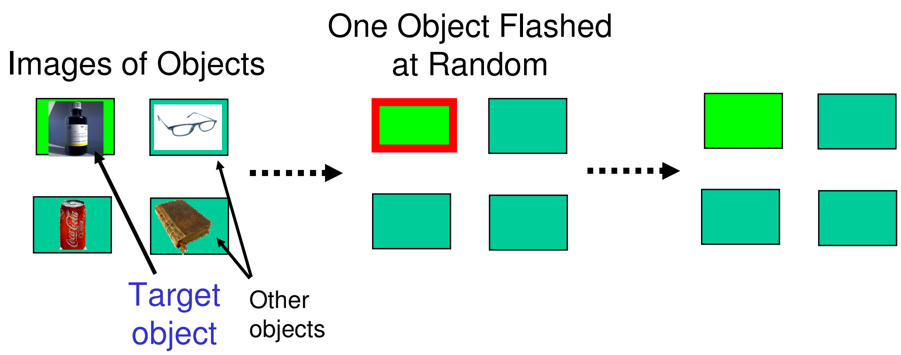

Brain response is different if flashed object is desired target object vs. non-target object

17

Example: Brain Responses Detected on Scalp

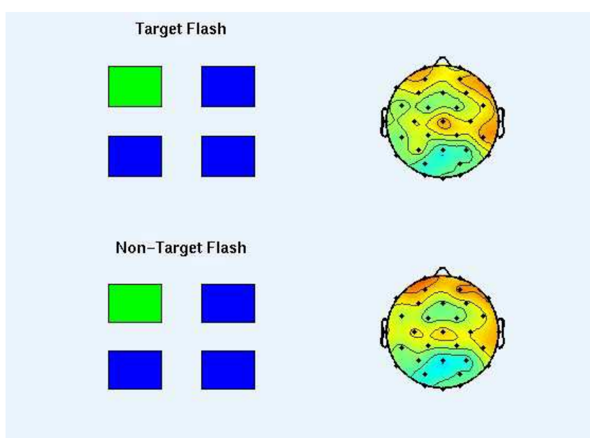

Use supervised learning to classify brain response as target or non-target

18

Example Task: Fetching an Object

Command: "Go to kitchen"

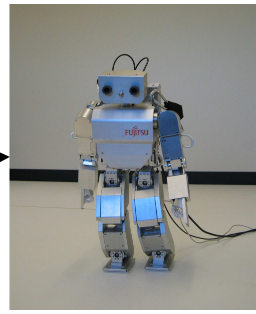

Robot navigates to kitchen

Command: "Pick object X in camera image"

Robot picks selected object

Command: "Bring object to me"

Robot brings selected object

19

Brain-Actuated Control of a Humanoid Robot

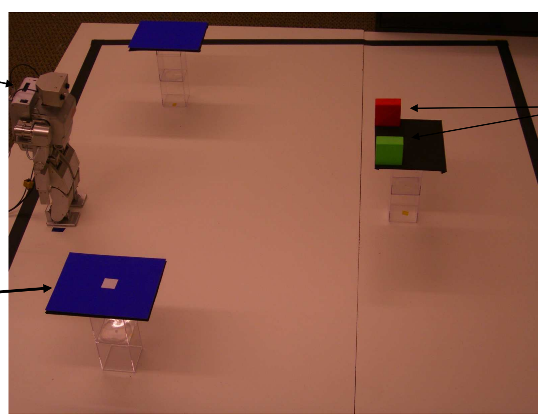

Task: Use brain signals to command robot to fetch red (or green) object and bring it to one of the tables.

(work of students Rawichote Chalodhorn, Ravi Kiran, Pradeep Shenoy, and CJ Bell) 20

Brain-Actuated Control of a Humanoid Robot

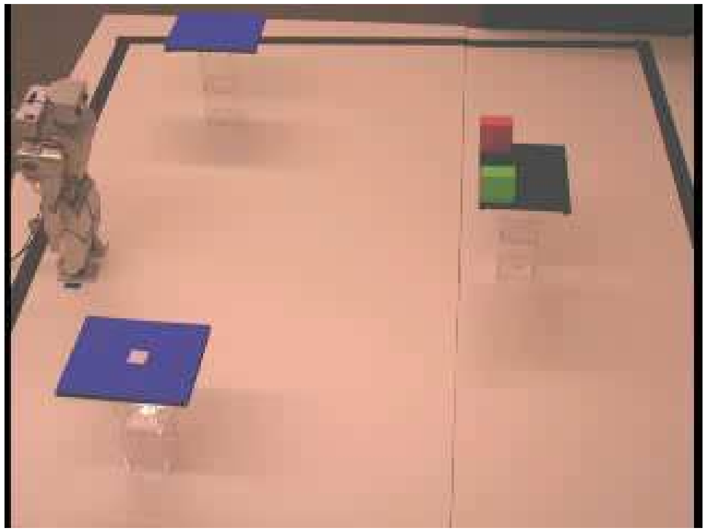

21

Take-Home Final: Details

- Will be posted on website later today
- 5 problems, open book, open notes
- Focus mostly on post-midterm material
- Due Wednesday Dec 13 by midnight via email to Raj and Abhay
- Will involve a mix of problem solving and descriptive questions

  E.g., Computing probabilities in Bayesian networks, explaining important concepts in AI (A* search, alpha-beta pruning, etc.)

22

Othello Tournament Update

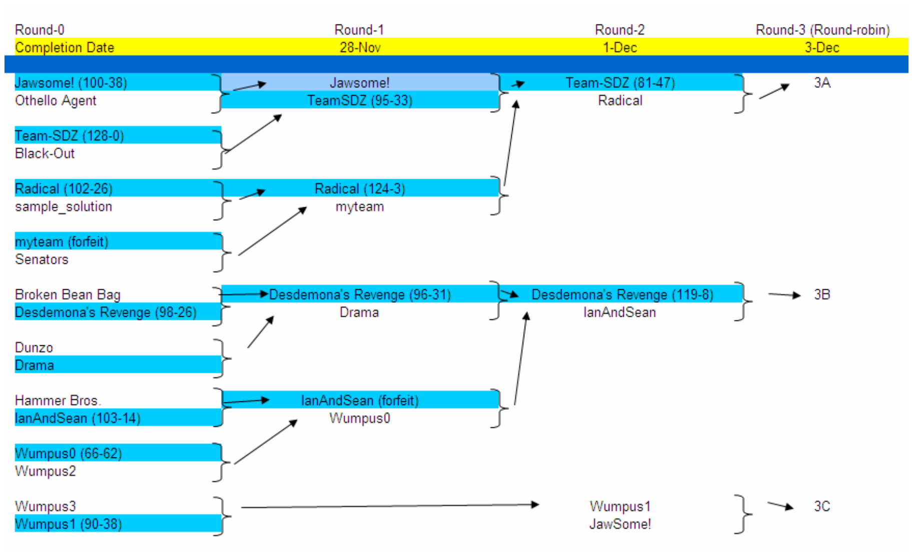

23

$$w_{kj} \rightarrow w_{kj} - \varepsilon \frac{dE}{dw_{kj}} \quad \text{But}: \frac{dE}{dw_{kj}} = \frac{dE}{dx_j} \cdot \frac{dx_j}{dw_{kj}}$$

$$\frac{dE}{dw_{kj}} = \left[-\sum_{m,i} (d_i^m - v_i^m) g'(\sum_j W_{ji} x_j^m) W_{ji} \right] \cdot \left[g'(\sum_k w_{kj} u_k^m) u_k^m \right]$$

24


Have a great break!
Revving Up the Romance, Passion & Excitement!
Who glued my fingers?