


CSEP 573
Final Lecture

Artificial Intelligence:
Past, Present, and Future



Plan for Today

- Part I
 - AI History and Review
 - Select Applications
 - The Future: where do we go from here?
- Break (student evals)
- Part II
 - Emerging Area in AI:
"Brain Computer Interfaces"
 - Guest lecture by Dr. Reinhold Scherer

A long time ago in a galaxy far away...

August 31, 1955

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College

M. L. Minsky, Harvard University

N. Rochester, I.B.M. Corporation

C.E. Shannon, Bell Telephone Laboratories

"We propose that a 2 month, 10 man study of artificial intelligence ...

An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.

We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together *for a summer.*"

Samuel's Checkers program (1959)

- First self-learning AI program
- Used *search tree* of board positions reachable from current position





Flashback: Search



- Uninformed Search

DFS and BFS

Iterative Deepening

Bidirectional Search

- Informed Search

Best first search:


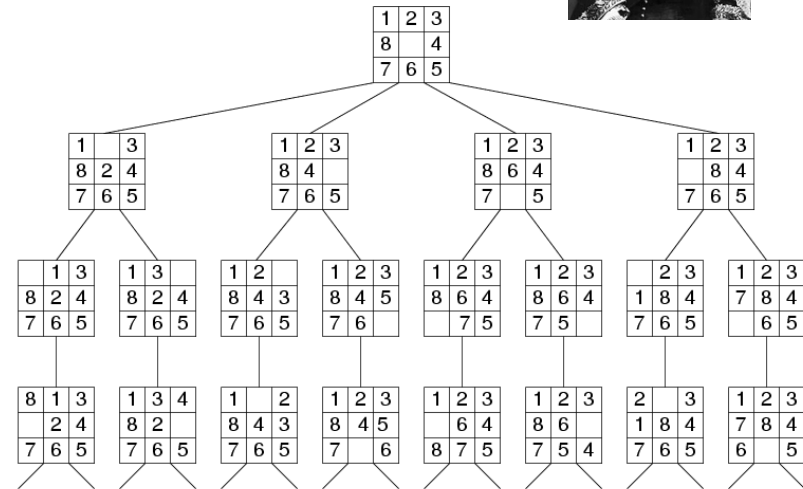
Greedy, A*, admissible heuristics

- Local Search

Hill Climbing

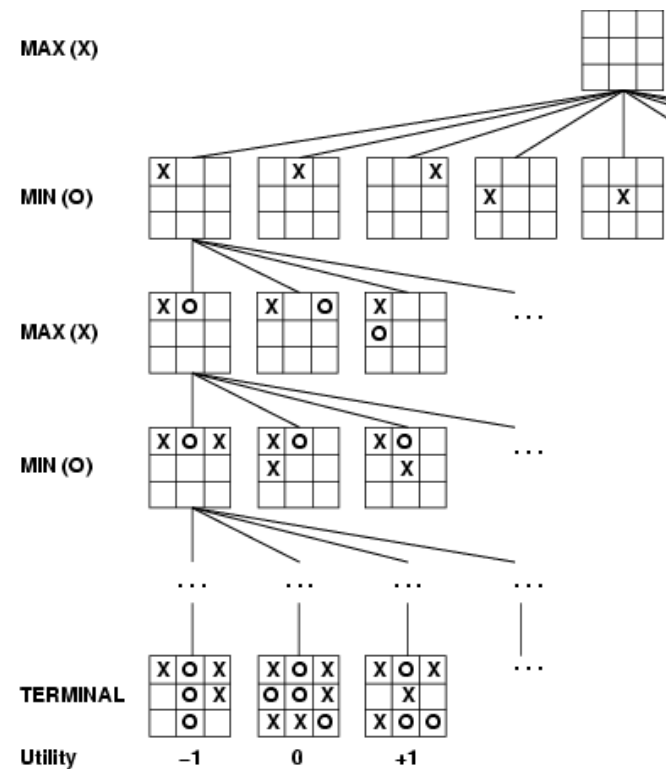

Simulated Annealing

Genetic Algorithms



Flashback: Adversarial Search

- Minimax Search
- Alpha Beta Pruning
- Truncated search and evaluation functions



Samuel's Checkers program (1959)

- First use of minimax search
- First use of alpha-beta pruning
- First use of truncated search and evaluation functions
- First use of machine learning
- Implemented on an IBM 701 with 9 KB memory!
- IBM's stock went up 15 points after demo



1961: First Industrial Robot




Unimate by Unimation

- Worked on a General Motors assembly line
- Transported die castings from assembly line and welded these onto auto bodies
- Began the era of industrial robots

Flashback: Robots today

- Invited Talks by:
- Dieter Fox (Probabilistic localization in robots)
- Rawichote Chalodhorn (Robot programming by human demonstration)



Math Flashback: Recursive Bayesian Updating


$$P(x_n | z_1, \dots, z_n) = \frac{P(z_n | x_n, z_1, \dots, z_{n-1}) P(x_n | z_1, \dots, z_{n-1})}{P(z_n | z_1, \dots, z_{n-1})}$$

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

$$\begin{aligned} P(x_n | z_1, \dots, z_n) &= \frac{P(z_n | x_n) P(x_n | z_1, \dots, z_{n-1})}{P(z_n | z_1, \dots, z_{n-1})} \\ &= \alpha P(z_n | x) \sum_{x_{n-1}} P(x_n | x_{n-1}) P(x_{n-1} | z_1, \dots, z_{n-1}) \end{aligned}$$

normalize *New data* *Dynamics* *Previous estimate*

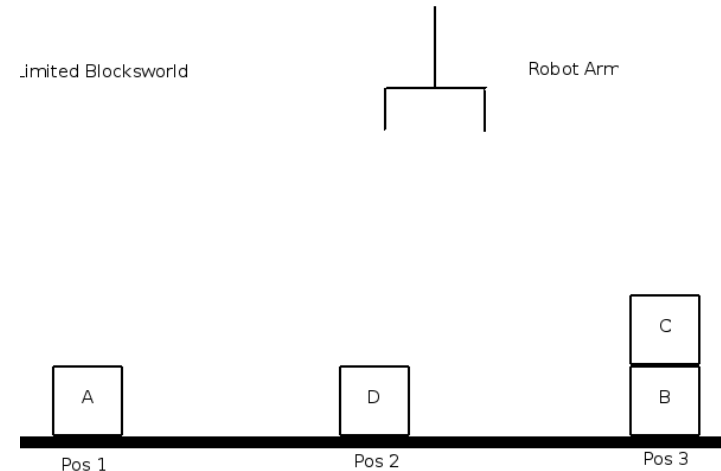
Application: Robot Localization and Mapping of Allen Center



(Work of Prof. Dieter Fox and students)

1971: Dawn of Classical Symbolic AI

- Blocks world model
Introduced by Terry
Winograd
- World is modeled as a set of
abstract symbols which may
be reasoned about using *logic*




Flashback: Logical Reasoning

- Propositional logic

Models and Entailment

Inference techniques:


- Soundness, completeness
- Resolution
- Forward/backward chaining
- WalkSAT



- First-Order Logic


Variables, Quantifiers Inference techniques:

- Skolemization & Unification
- Forward/backward chaining
- Resolution




1980s: Neural Networks

- **Early neural networks**
 - McCulloch & Pitts (1943) – simple neural nets
 - Rosenblatt (1962) – perceptron
- **Backpropagation learning algorithm**
 - Invented in 1969 and again in 1974
 - Hardware too slow, until rediscovered in 1985



Flashback: Neural Networks



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

Backprop rule for input-hidden weights w :

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

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

Application: Handwriting Recognition

Artificial Neural Network Handwriting Recognizer

Written in Java™

by [Bob Mitchell](#)



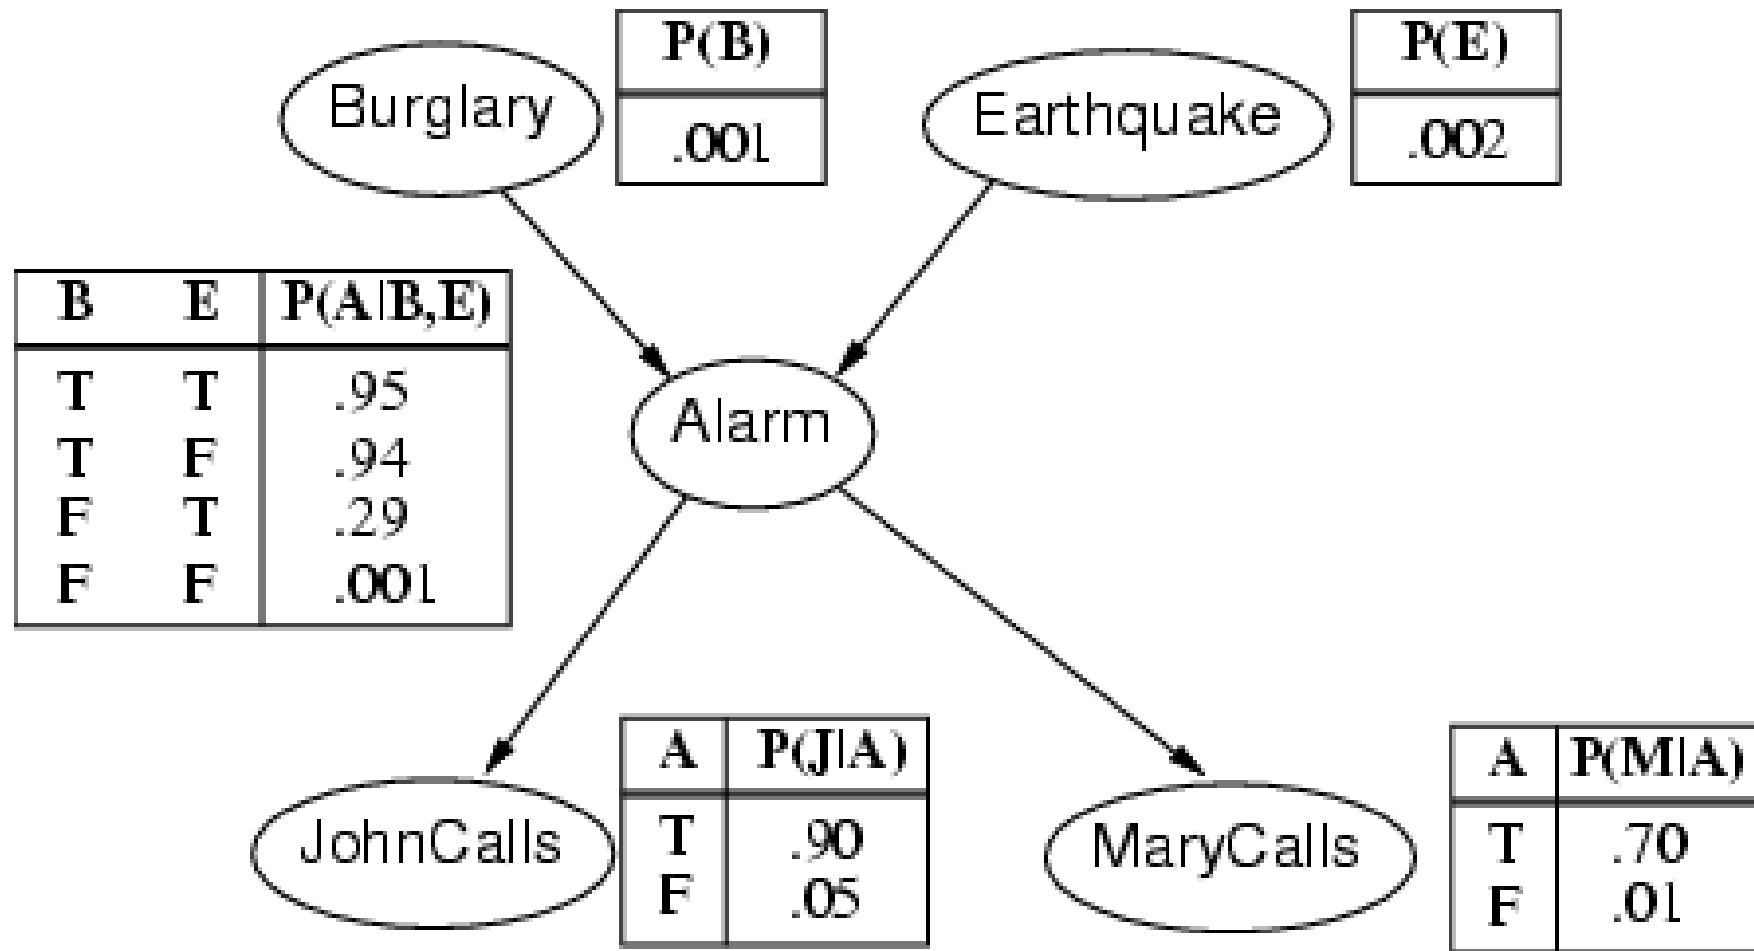
[Demo](#)

<http://www.cs.rochester.edu/~kautz/Courses/290Bspring2008/NeuralNets/NeuralNetsHandwriting/JRec.html>

1990s to present: Probabilistic Models and Machine Learning

- Probabilistic graphical models
 - Pearl (1988) – Bayesian networks
- Machine learning
 - Quinlan (1993) – decision trees (C4.5)
 - Vapnik (1992) – Support vector machines (SVMs)
 - Schapire (1996) – Boosting
 - Neal (1996) – Gaussian processes
- Recent progress:
 - Probabilistic relational models, deep networks, active learning, structured prediction, etc.

Flashback: Bayesian Networks




Application: Tracking a Person using GPS

Tracking



Goal Prediction




Anomaly Detection



foot=blue, bus=green, car=red

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


Application: Robot Learning by Imitation




Imitating from Motion Capture Data



Motion Capture



Data from Motion Capture




Attempted Imitation

Bayesian Network for Stable Imitation and Learning

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


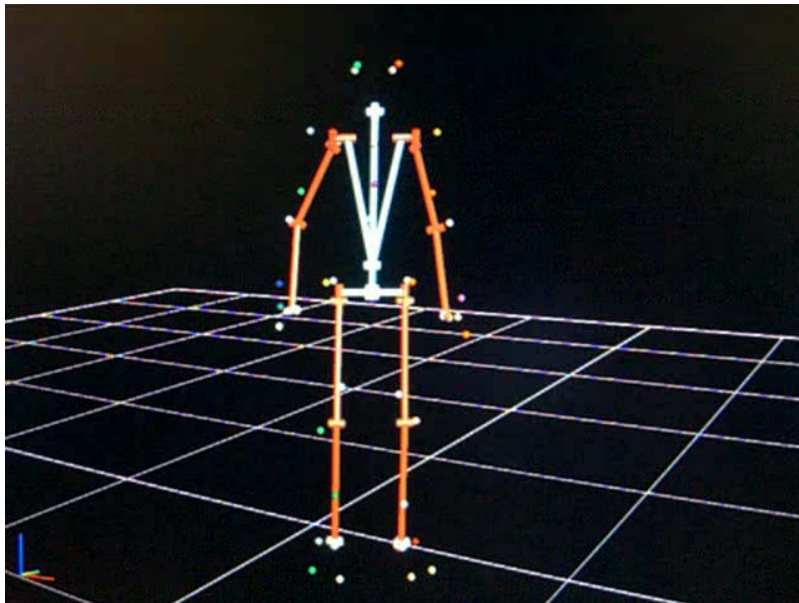
State \mathbf{s} = [joint angles, gyro values, foot pressure values]

Action \mathbf{a} = [position commands to motors for each joint]




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


Learning to Imitate a Human Action



Result after Learning



Human Action



Imitation

The Future of AI

Massive amounts of data

+

Sophisticated probabilistic reasoning
and machine learning algorithms

+

Massive computing power

= AI revolution?

Automated Driving



Winners of the 2005 and 2007 DARPA Grand Challenges

Driverless pod cars at Heathrow International Airport



AI in a Sensor-rich World

- Intelligent houses
- Intelligent refrigerators
- Intelligent forests
- Intelligent oceans
- Intelligent bridges
- Etc.

AI in Industry

- **Joseph Sirosch's talk:**
 - Fraud detection, trust, and safety
 - Just-in-time inventory systems
 - Collaborative filtering
 - Recommendation in social networks
 - Behavioral ad targeting
- **Other applications**
 - Stock market prediction
 - Insider trading and market abuse detection
 - AI-assisted design
 - Intelligent robots for manufacturing and testing

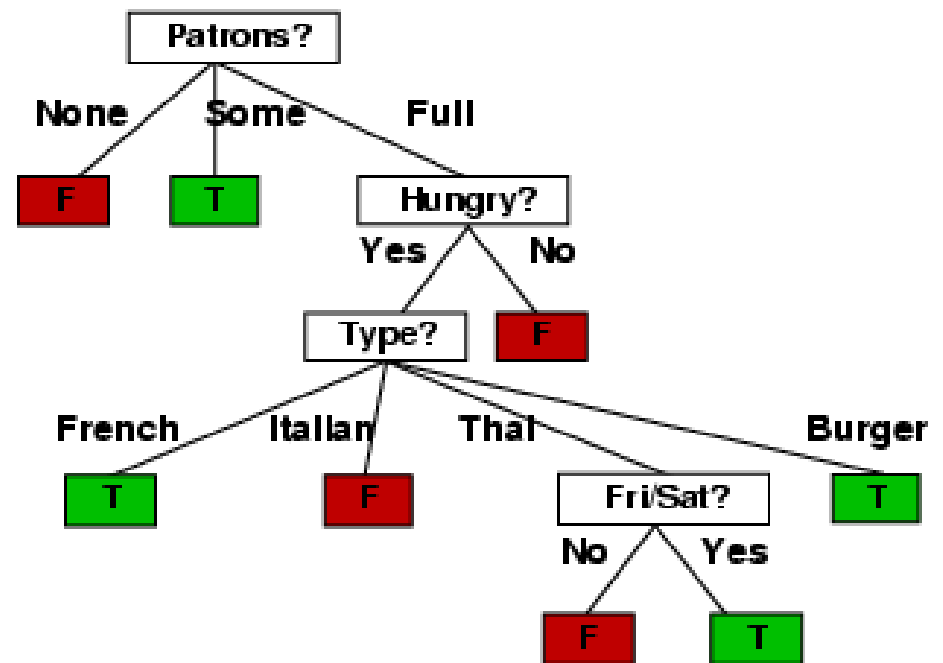
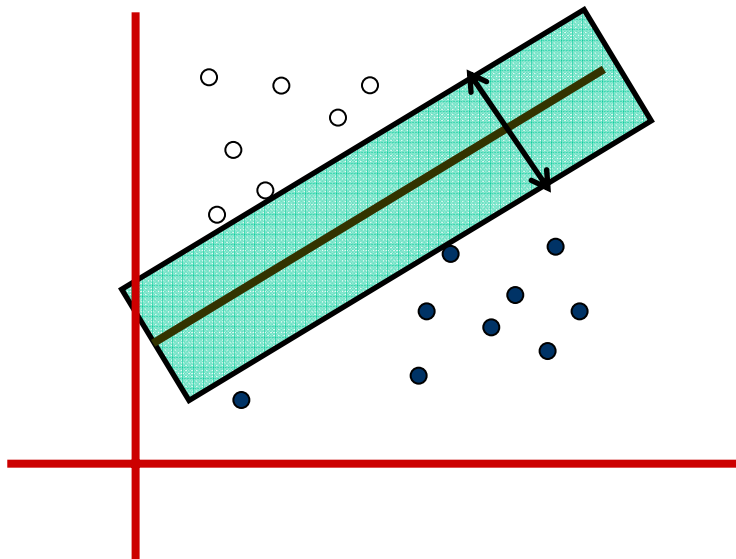
Other future AI applications

- **Smart power grids:** electric power flows both ways and is distributed dynamically according to changing demand
- **Security and military:** Bomb diffusing robots, unmanned vehicles, "soldier" robots
- **Robot firefighters**
- **AI Travel Agents**
- **AI Accountants**
- **AI Cashiers**
- **AI Football Coaches**
- **AI Football Players**
- ...

Flashback: Machine Learning


Classification

- Decision Trees
- Neural Networks
- SVMs




Application: Brain-Computer Interfaces


- Classifying brain signals recorded at the scalp
- Detect what a person wants from a set of options
- Command a humanoid robot to fetch an object
- Details in Dr. Scherer's talk



[CBS News Sunday Morning](#)



Thank you for your attention!



10 minute break

