


## The Games Computers

### (and People) Play

**Jonathan Schaeffer**  
Department of Computing Science  
University of Alberta  
Edmonton, Alberta  
Canada T6G 2E8  
jonathan@cs.ualberta.ca

## The Computer (R)Evolution


Need to re-think what it means to think.



## AI and History

One of the most profound contributions of the 20<sup>th</sup> century is the realization that intelligent behavior can be realized by non-human information processing architectures.


## Different Architectures...



## ... Different Solving Methods



How?  
Why?  
What?  
Where?  
When?



## Looking Under the Hood...




Peek inside a "smart" program and you may find little that is "intelligent". Yet computers do "intelligent" things.


Is this a paradox?  
What is intelligence?


AI creates the illusion of intelligence

## ... Looking Under the Hood



Of course, some humans try to create the illusion of intelligence too!

Disadvantage of the human architecture: 

An architecture comparison: 

## Exploit the Strengths

- Brain architecture
  - ♦ language and vision processing
  - ♦ generalization
  - ♦ reasoning by analogy
- Computer architecture
  - ♦ calculations & exhaustive enumeration
  - ♦ repetitious tasks
  - ♦ large, infallible memory

## Avoid the Weaknesses

- Brain architecture
  - ♦ calculations & exhaustive enumeration
  - ♦ repetitious tasks
  - ♦ large, infallible memory
- Computer architecture
  - ♦ language and vision processing
  - ♦ generalization
  - ♦ reasoning by analogy

## Research and Games

- Games have been a popular choice for experimenting with artificial intelligence techniques, mathematical study and economic models.
- Success in building game-playing programs has resulted from exploiting the strengths of the computer, and down-playing the human example and the mathematical literature.


## Progress in Computer Games

- Survey the astounding success in building programs capable of challenging human supremacy
- Bring a human side to Man versus Machine: The Experiment
- Discuss the computer's secrets of success

## (1) Search

- Simple algorithmic solution for computing the game-theoretic value for two-person, zero-sum games...
- ... but impractical for large search spaces
- Combine deep online and offline searches to approximate the game-theoretic value of a position

## Checkers



## Chinook

Have to overcome the stigma of checkers being "solved" in 1963.

Project takes five years, 10 people, > 200 computers working around the clock, and terabytes of data.



## Human Perfection!?



Name: Marion Tinsley  
Profession: Teach mathematics.  
Hobby: Checkers  
Record: Over 42 years loses only 3 (!) games of checkers.  
Crime: Too good at checkers.

## The Verdict...

- The first computer to win a human world championship!
- Computers will never be overtaken by man at checkers.




## The Secret?

- Endgame databases
  - ♦ searched all positions with 8 or fewer pieces
  - ♦ each identified with perfect win, loss, draw info
  - ♦ 444 billion positions in the program's memory
  - ♦ exceeds human abilities
  - ♦ introduces perfect knowledge into the search
  - ♦ factual knowledge, but without the ability to generalize it


## Fallible Humans

The 100-Year Position  
Give it to humans for 100 years... win!



Give it to Chinook for 5 seconds... draw!  
The 197-Year Position

## Chess



## Man Versus Machine

<u>Kasparov</u>	<u>Name</u>	<u>Deep Blue</u>
5'10"	<b>Height</b>	6' 5"
176 lbs	<b>Weight</b>	2,400 lbs
34 years	<b>Age</b>	0.5 years
50 billion neurons	<b>Computers</b>	512 processors
2 pos/s	<b>Speed</b>	200,000,000 pos/s
Extensive	<b>Knowledge</b>	Primitive
Electrical/chemical	<b>Power Source</b>	Electrical
Enormous	<b>Ego</b>	None

## Bravery and Bravado




Name: Garry Kasparov  
Title: World Chess  
Champion.  
Crime: Over confidence?

What went wrong?  
Who is the better chess player?



## The Verdict...

- Exhibition match.
- Scientific data point can't be repeated.
- Man was superior in 1997 but...
- Junior/Fritz matches give new insights




## The Secret?

- Brute-force search
  - ♦ consider all moves as deeply as possible
  - ♦ some moves can be provably eliminated
  - ♦ 200,000,000 *per second* versus Kasparov's 2
  - ♦ 99.99% of the positions examined are silly by human standards
  - ♦ lots of search... and little knowledge
- *Tour de force* for engineering

## (2) Knowledge... Sort Of

- Let computers discover their own knowledge about a domain
- Training with numerous examples and self-play allows the program to discover useful knowledge and tune it

## Backgammon



## TD-Gammon



Had to overcome the stigma of backgammon being "solved" in 1979.

Gerry Tesauro builds TDGammon over 8 years. It can learn to play strong backgammon.

## Fearless Fighter



Name: Malcolm Davis  
Title: World backgammon champion.

Crime: Agrees to play exhibition matches against a computer.

Consequence: Narrowly avoids becoming part of computing history.

## The Verdict...


World champion caliber but hard to prove it is the world champion.



## The Secret

- Pioneering success for temporal difference learning
- Combination of search, expert knowledge, and a neural net tuned using TD learning
- *Tour de force* for artificial intelligence

## Reversi/Othello



## Logistello



Had to overcome the stigma of Othello being "solved" in 1980 and 1990.

Michael Buro's one-man effort for five years produces Logistello.

## Cold Blooded Murder



Name: Takeshi Murakami  
Title: World Othello Champion  
Crime: Man crushed by machine.



## The Verdict...


Man will never be close to machine again.




## The Secret?

- Automatically discovered and tuned knowledge
  - samples patterns to see if its presence in a position can be correlated with success
  - identified 11 patterns (46 with rotations)
  - tuned 1.5 million parameters using self-play games with feedback
- “Knowledgeable” program but no one understands the knowledge

## Crossword Puzzles



## Proverb – Duke University



## The Secret?

- Multiple agents
  - ♦ 34% of clues/answers are repeats
  - ♦ specialized solvers (puns, word plays)
  - ♦ specialized searches (dictionaries; geography, history, medical, movie databases; etc.)
- Proverb scores well on the NY Times puzzles *without understanding the clues*
- Information retrieval without understanding the information is a powerful technique!

## (3) Statistical Sampling

- Search for games of non-deterministic and/or imperfect information
- Sample from the space of possibilities to approximate the true value

## Bridge



## Why is this Man Flying High?

- Matt Ginsberg develops the first expert-level bridge program, GIB (1998).



- Finishes 12th in the World Championship.

## When the Going Gets Tough...



Name: Zia Mahmood  
Motive: In 1990 offers  
£1,000,000 bet that  
no program can defeat  
him.  
Crime: December 1, 1996  
cancels bet when faced  
with a possible  
challenger.

## The Verdict...

- Man is better than machine!
- But... Ginsberg may be only a few years away from success



## The Secret?

- GIB does 100 simulations for each decision
  - ♦ deals cards to opponents consistent with available information
  - ♦ chooses the action that leads to the highest expected return
  - ♦ program does not understand things like “finesse” or “squeeze”
  - ♦ simulations contain implicit knowledge

## Poker



## Poki – University of Alberta



Have to overcome the stigma of poker being “solved” in 1985.

Poker is a hard problem because of multiple opponents, imperfect information, and deception


## Know When to Hold'em...



... and if you are 3-time World Champion Stu Ungar, you don't want to fold'em just yet!

## The Verdict...

- Poker is safe for a few years to come.
- Quickly adapting to changing opponents is a challenging problem.



## The Secret?

- Seeds of human defeat are present...
- Precise probability calculations
- Game theoretic solutions (Koller & Pfeffer)
- Simulations to approximate probability distributions
- Use short-term and long-term statistics to model each opponent

## Rock, Paper, Scissors

- 1st & 2nd International RoShamBo Championships
- Deceptively simple problem that illustrates the differences between optimal and maximal play



## locaine Powder



## The Priorities of Life

- Theorem (F. Gump):
  - ♦ “Life is a box of chocolates”

Corollary (M. Ginsberg)

“Winning a box of chocolates is life.”



Bet: Ginsberg can't achieve his goal by 2003

Bet: Schaeffer can't achieve his goal by 2003

## Scrabble



## Maven




Brian Sheppard spends 14 years developing his Scrabble program.

## Should this Man Be Happy?



Name: Adam Logan  
Profession: Math professor.  
Qualifications: 1997 Canadian and North American champion.

## Maven versus Logan: A Classic



## The Verdict...

Computers are better than humans and the gap will only grow with faster computers




## The Secret?

- Memory
  - ♦ Maven has the entire dictionary in its memory
  - ♦ over 100,000 words
- Simulations
  - ♦ simulates 1,000 game scenarios per decision
  - ♦ typically 700 legal moves (more with a blank)!
  - ♦ becomes a constraint-satisfaction and optimization problem

## (4) Help!

- Some games are resistant to all known techniques
  - ♦ mathematics is of little help
  - ♦ too hard for search alone
  - ♦ knowledge is inadequate
- Hex, amazons, octi, and



## Go



## Why is this Man Pensive?



Name: Chen Zhixing  
Author: Handtalk (Goemate)  
Profession: Retired  
Computer skills: self-taught assembly language programmer.  
Accomplishments: dominated computer go for 4 years.

## Why is this Boy Happy?



Because he can give Goemate a 9 stone handicap and still easily beat the program, thereby winning \$15,000.

## The Verdict...

\$1,000,000 prize is safe for many decades to come.



## The Problem?

- Brute-force search will not work
  - ♦ the only approaches we know of involve extensive knowledge
  - ♦ roughly 60 major knowledge-based components needed
  - ♦ program is only as good as the weakest link
- We have no idea how to tackle this domain effectively with computers

## Challenges!

- Explaining computer calculations
  - ♦ have computers write the definitive books (Frank)
- Annotating games
  - ♦ annual ICCA competition
- Extending learning
  - ♦ knowledge, search control, algorithm
- Feature discovery
  - ♦ data mining looks promising



## Commercial Games

- Hard problems!
- Real-time, limited memory, limited time, must work, must provide complete coverage



## Perspective on Games: Pro

“Saying Deep Blue doesn’t really think about chess is like saying an airplane doesn’t really fly because it doesn’t flap its wings”

Drew McDermott

## Perspective on Games: Con

“Chess is the *Drosophila* of artificial intelligence. However, computer chess has developed much as genetics might have if the geneticists had concentrated their efforts starting in 1910 on breeding racing *Drosophila*. We would have some science, but mainly we would have very fast fruit flies.”

John McCarthy

## The Reality

- Most of these projects took tens of man years to solve “easy” problems.
- What does this say about the difficulty of building high-performance solutions to important real-world problems?

## For More Information

- Special issue of *Artificial Intelligence*
  - ♦ January, 2002
- *Chips Challenging Champions*
  - ♦ Schaeffer and van den Herik (editors)

## University of Alberta

- Host for AAAI National Conference
- Host for Computers and Games 2002
  - ♦ [www.cs.ualberta.ca/cg2002](http://www.cs.ualberta.ca/cg2002)
- Looking for students interested in...
  - ♦ AI using games as an experimental test-bed
  - ♦ board and card games
  - ♦ Electronic Arts and BioWare products