


Artificial Intelligence



TSP Competition

1. Xue Liu '12. 16247.4804.
2. Wesley Verne '13. 16314.1572.
3. Sebastian Steffen and
Linda Zhao '14. 17003.623.



Smallest insertion. 17265.628.

Nearest insertion. 27868.710.

Best so far. Ilias Giechaskiel '13. 16154.548

Artificial Intelligence

Fundamental questions.

- Is real life described by discrete rules, or not?
- Can we build a UTM from living components?
- Can a machine do anything a human can do?
- Can human intelligence be simulated by a machine?

Artificial Intelligence (AI). The science and engineering of making intelligent machines.



human brain
(around 100 billion neurons)



Intel Core i7 microprocessor
(around 781 million transistors)


Weak AI

Can a machine appear intelligent?

Artificial Intelligence

Goal. [Turing 1950] Program computer to exhibit intelligent behavior.


*“ Every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. ” — J. McCarthy *51*



1960s. Very optimistic predictions.


Reality. Slow progress with some striking successes.

The Chess-Playing Turk



Tic-Tac-Toe

Tic tac toe. Two person game of skill.




Number of possible games. 255,168.

Tic-Tac-Toe

Minimax algorithm.

- If game over, return winner.
- If X's turn, return max
- score of children.
- If Y's turn, return min
- score of children.



Tic-tac-toe is solved. X will never lose; O will never lose.

assuming optimal play by X

Chess

Challenge. [Claude Shannon] Develop a computer program to play chess.



Number of possible games. $\approx 10^{23}$.

N-by-N version. EXPTIME-complete.

Kasparaov vs. Deep Blue

Deep Blue. [IBM]

- Supercomputer, augmented by VLSI chess chips.
- 200 million board positions per second.




Machine beats man. [February 1996] First computer program to win a chess game against reigning world champion.

Q. Does Deep Blue appear intelligent?

modern chess programs run on a laptop are now even better than Deep Blue

Checkers (Draughts)




Number of possible games. $\approx 10^{31}$.
N-by-N version. EXPTIME-complete.

Checkers (Draughts)

Chinook. [Jon Schaeffer] Computer program for checkers.

Man vs. machine. Chinook awarded world championship in 1994 after 6 draws with Marion Tinsley (who withdrew).



Checkers is solved! [Science, September 2007]

- Proof that black will never lose with optimal play.
- Proof that white will never lose with optimal play.

Backgammon

Backgammon. Two-player game of skill and luck.



TD gammon. [Gerry Tesauro 1980s]

- Program was given no expert backgammon knowledge.
- Learned strategy by playing itself 300,000 times.
- Among best backgammon players in the world.

Poker

Poker. Multi-player game of skill and luck.



Texas hold 'em poker bots. Not currently competitive with pros.


Turing Test

Can a machine appear intelligent?
If so, how could we tell?

Turing Test

Judge types questions to both.

- Computer claims to be human.
- So does human.



Pass. If judge cannot tell which entity is human, we say computer has passed the test and is "intelligent."

Prediction. [Turing] Computer would pass test by 2000.

Eliza

Eliza. [Weizenbaum 1966]

- Program imitates Rogerian psychiatrist.
- Works by generating scripted responses to keywords.
- Just 200 lines of code!

Q. Does Eliza appear intelligent?

A. You be the judge: <http://chayden.net/eliza/Eliza.html>



Eliza Pattern Matching Rules

...

key: perhaps

decomp: *

reasm: You don't seem quite certain.

reasm: Why the uncertain tone ?

reasm: Can't you be more positive ?

reasm: You aren't sure ?

reasm: Don't you know ?

key: francais

decomp: *

reasm: goto xforeign

reasm: I told you before, I don't understand French.

key: sorry

decomp: *

reasm: Please don't apologise.

reasm: Apologies are not necessary.

reasm: I've told you that apologies are not required.

key: apologise

decomp: *

reasm: goto sorry

...

Parry

Parry. [Colby 1972] Program imitates a paranoid schizophrenic.

Q. What happens when Parry "meets" Eliza?

A. You be the judge: <http://tools.ietf.org/html/rfc439>

Turing Test

Loebner Prize

- \$100,000 to first program to pass restricted Turing Test.
- Annual competition held since 1995.
- Small awards given to best programs.



Loebner Contest Dialogs

“Hello, Judge. My name is Ella. Please tell me your first name, or whatever I should call you.”

“Call me Lisa. Am I male or female?”

“ Okay, I will call you 'Lisa. Am I Male Or Female?' I am here recently from Tainjin, China, so please be patient with me.”

“ I'm not a patient because I am not ill. Are you feeling well today?

“ What do you think it means that you are not ill?”

“ Have you ever been ill?”

“ 'Have you ever been ill?' you say? That's it, Lisa Am I Male Or Female! I will schedule a whipping for you.”

Loebner Contest Dialogs

"What do you think of Reading?"

"I'm somewhat crazy about reading. I like to read mystery novels."

"I meant the town."

"Jones lives in the city and is a writer of children's books."

"Does Jones live in Reading, then?"

"Tomorrow is another day. We can only live today well."

"Do you live, then?"

"I live then a lot. Someday we will live on another planet."

"I think you are already there."

Loebner Contest Dialogs

"What is it?"

"What is what?"

"A high-yield, perennial, four-lobed hybrid of wheat and rye whose root grain, triticale, can trace its ancestry to twentieth-century Canada."

"At last. Another classic grain freak."

"Lewis. Jim Lewis. Would you like for me to show you some poetry?"

"No, thanks."

"I didn't want to show you the stupid poetry anyway."

Digression: "Reverse" Turing Test

Standard Turing test. Judge is human.

Reverse Turing test. Judge is computer!

Why?



- Google allows each user 7GB storage.
- PayPal once offered \$5 for each user who opens a new account.
- Both need to distinguish real humans from bots.



Exploiting Intractability: Captcha's

OCR. Given degraded text, find original text.

CAPTCHA. [completely automated public Turing test to tell computers and humans apart]



http://online.wsj.com/public/resources/images/OB-AB313_captch_20060524170113.gif

Knowledge Databases

Twenty questions. <http://www.20q.net>

Question answering system. <http://start.csail.mit.edu>



Q. Does a computer than can answer questions appear intelligent?

IBM Watson

IBM Watson. Capable of answering questions posed in natural language.

- Combination of natural language processing, information retrieval, machine learning, and massively parallel computing.
- Massively parallel computer.
- Written in Java and C++.

Challenge. Play Jeopardy competitively against best human players.



Watson vs. Brad Rudder vs. Ken Jennings



Will Watson Put Your Job in Jeopardy?

Q. What to do with Watson now?

A. Many commercial applications in the works.

- Business logic.
- Medical diagnoses.
- Legal research.
- ...


“The goal is to have computers start to interact in natural human terms across a range of applications and processes, understanding the questions that humans ask and providing answers that humans can understand and justify. ” — The DeepQA Project

Will Watson Put Your Job in Jeopardy?

Q. What to do with Watson now?

A. Many commercial applications in the works.

- Business logic.
- Medical diagnoses.
- Legal research.
- ...



Strong AI

*Can a machine **be** intelligent?*

Chinese Room Experiment (Searle 1980)

Imagine that:

- You don't understand Chinese.
- You're alone in a room that has paper slots labeled "input" and "output."
- You have a big book of Chinese writing.
- You have English instructions (**but no translations**) that tell you what to write on your output paper in response to various inputs.



Chinese Room Experiment (Searle 1980)

And then:

- Chinese speakers outside the room pass in pieces of paper with Chinese writing. They know these are questions (but you don't).
- You consult your manual of instructions, figure out the proper Chinese response, copy it down, and pass it out.

什麼帶來
快樂
爲天下式

If you see this shape,
"什麼"
followed by this shape,
"帶來"
followed by this shape,
"快樂"

then produce this shape,
"爲天"
followed by this shape,
"下式".



Chinese Room Experiment (Searle 1980)

- Q. The folks outside think you understand Chinese. Do you?
- Q. If a computer did the same, would it understand Chinese?



Chinese Room Experiment

Weak AI. Can machines be programmed to **exhibit** intelligent behavior?

A. Surely true: Deep Blue, Chinook, TD-Gammon, Watson, ...

Strong AI. Can machines can be programmed to **possess** intelligence?

Searle. Chinese Room is absolute refutation of strong AI.

But... many disagree!

“ The question of whether a computer can think is no more interesting than the question of whether a submarine can swim. ” – Edsger Dijkstra

An interesting question, but... either way, limitless possibilities remain for applications of computer science.

DARPA Grand Challenge

DARPA Grand Challenge

2004 Grand Challenge. Navigate an autonomous vehicle through 142 mile course in Mohave Desert at military speed.

Results. No team finished; CMU team finished 7.36 miles.



*Prospect Eleven
(a Princeton undergraduate team!)*

DARPA Grand Challenge

2005 Grand Challenge. Navigate an autonomous vehicle through 132 mile course in Mohave Desert at military speed.

Results. Stanford team won in under 7 hours; \$2 million prize.



Prospect Eleven
(successfully navigated 9.4 miles)



Stanley
(Stanford Racing Team)

failed because of a performance bug,
which starved code to control steering

DARPA Grand Challenge

2007 Urban Challenge. Navigate an autonomous vehicle through 60 mile course in **mock urban** environment, while obeying traffic laws and avoiding other vehicles.



DARPA Grand Challenge

2007 Urban Challenge. Navigate an autonomous vehicle through 60 mile course in **mock urban** environment, while obeying traffic laws and avoiding other vehicles.



the six finishers of 2007 Urban Challenge

Overview (revisited)

What was COS 126? Broad, but technical, intro to computer science.


Goals.

- Demystify computer systems.
- Empower you to exploit available technology.
- Build awareness of substantial intellectual underpinnings.

Topics.

- Programming in Java.
- Machine architecture.
- Theory of computation.
- Applications to science, engineering, and commercial computing.

“ Computers are incredibly fast, accurate, and stupid; humans are incredibly slow, inaccurate, and brilliant; together they are powerful beyond imagination. ” – Albert Einstein





*“ Do not do what I do; rather, take whatever I have to offer
and do with it what I could never imagine doing and then
come back and tell me about it. ”* – Mark C. Taylor

Bye, congratulations, and good luck!