


Artificial Intelligence



Conway's Game of Life



John Conway



hacker's emblem

Conway's Game of Life


Conway's game of life.

- Critters live and die in an infinite square grid.
- Time proceeds in discrete steps.

Survival. Critter lives and dies depending on 8 neighbors:


- Too few? (0-1) die of loneliness
- just right? (2-3) survive to next generation
- too many? (4-8) die of overcrowding

Birth. Critter born if exactly 3 neighbors.




Conway's Game of Life


Glider. Propagates a signal.




time t




time t+1



time t+2




time t+3



time t+5

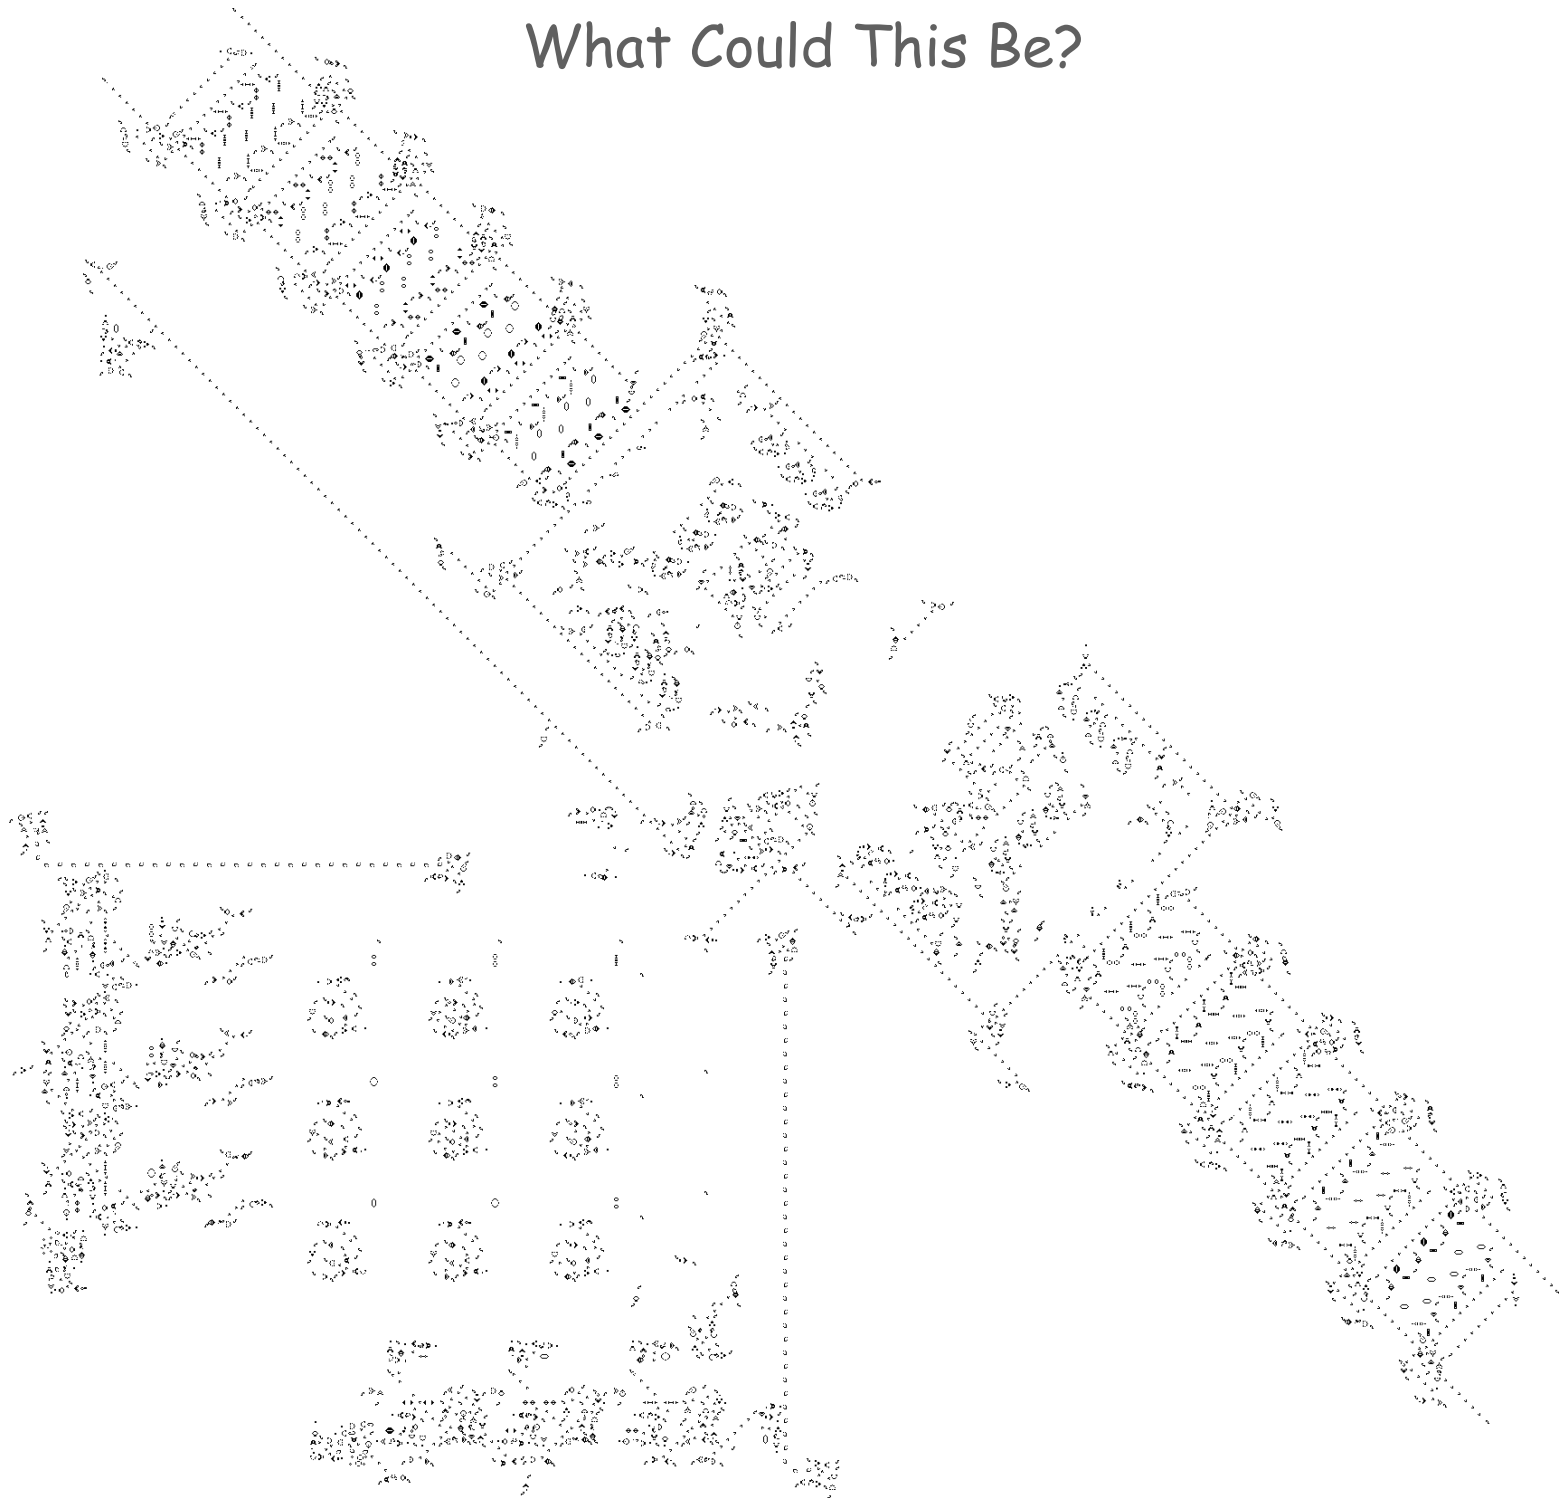
Conway's Game of Life

Gosper glider gun. Generates gliders.



<http://www.ibiblio.org/lifepatterns/>

What Could This Be?



Implications

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


Weak AI

*Can a machine **appear** to be 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.

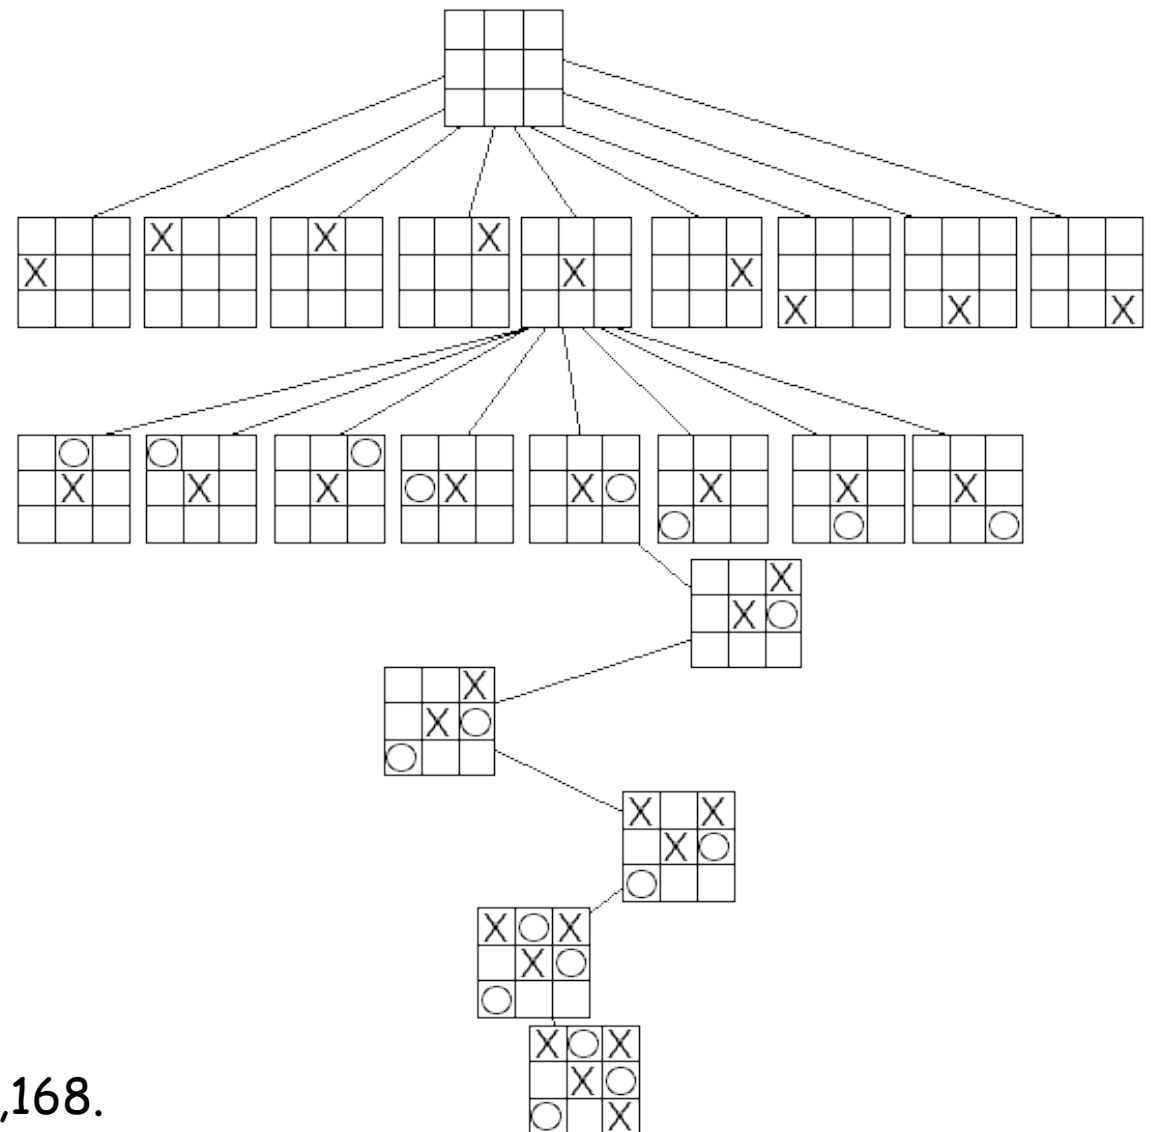
Reality. Slow progress with some striking successes.

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.




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. $\approx 10^{23}$.

Kasparov 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 a chess-playing computer program appear intelligent?

Chess



“ Once again, man beats machine! ”

Checkers (Draughts)




Number of possible games. $\approx 10^{31}$.

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.



Poker bots. Not currently competitive with pros.

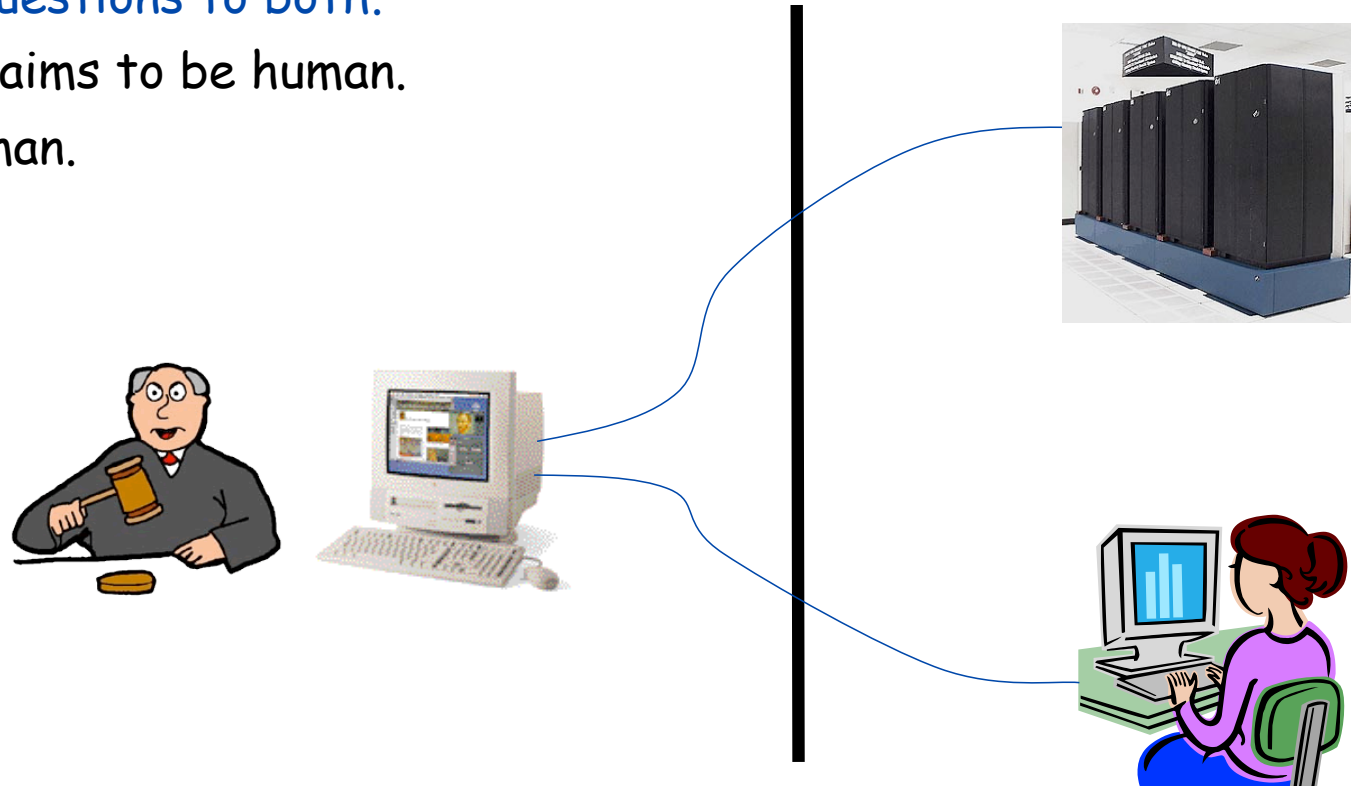
Turing Test

*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: *

reasmb: You don't seem quite certain.

reasmb: Why the uncertain tone ?

reasmb: Can't you be more positive ?

reasmb: You aren't sure ?

reasmb: Don't you know ?

key: francais

decomp: *

reasmb: goto xforeign

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

key: sorry

decomp: *

reasmb: Please don't apologise.

reasmb: Apologies are not necessary.

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

key: apologise

decomp: *

reasmb: 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 that can answer questions appear intelligent?

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



Princeton team

First DARPA Grand Challenge

Red Team. Reached mile 7.4. Vehicle went off course, got caught on an obstacle and rubber on the front wheels caught fire.

SciAutonicsII. Reached mile 6.7. Vehicle went into an embankment and became stuck.

Caltech. Reached mile 1.3. Vehicle went through a fence, and couldn't come back through.

CIMAR. Reached mile 0.45. Vehicle ran into some wire and got totally wrapped up in it.

Virginia Tech. Vehicle brakes locked up in the start area.

ENSCO. Vehicle flipped in the start area, experienced a fuel leak, and the team needed to shut off the fuel.

DARPA Grand Challenge

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



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 (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, others.

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

Searle. Chinese Room is absolute refutation of strong AI.

But... many disagree!

Longstanding debate: Is (Strong) AI Ultimately Possible?

Patrick Winston. "Just as the Wright brothers at Kitty Hawk in 1903 were on the right track to the 747, so too is AI, with its attempts to formalize commonsense understanding, on its way to fully intelligent machines."

Hubert Dreyfus. "Believing that writing these types of programs will bring us closer to real AI is like believing that someone climbing a tree is making progress toward reaching the moon."

Marvin Minsky. "The brain happens to be a meat machine."

Herb Simon. "Either AI is possible ... or we're not."

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

An interesting question, but...

either way, limitless possibilities remain for applications of computer science

A final thought

Shirley Tighlman. "[Princeton's Neuroscience program] has defined as its mission the **understanding of how the brain works**, how the wiring diagram is assembled during development, and how that wiring diagram results in complex processes like learning and memory and complex thinking..."

If we can understand it, can we **simulate** it?

Note: 100 billion neurons in the brain.
500 billion bytes in your PC.

Overview (revisited)

What is COS 126? Broad, but technical, introduction 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





Congratulations, and good luck!