

# Is Artificial Intelligence Possible?

or

## Can Machines Think?

Reasons for answering “No”.

- **Theology:** Machines don't have the equipment, *i.e.* no souls.
- **Consequences:** It would just be too awful if they could.
- **Logic:** There are formal proofs that there are some things that machines can never do.
- **Consciousness:** Machines aren't conscious, so they can't possibly think.
- **Determinacy:** Computers can only do what they are programmed to do.

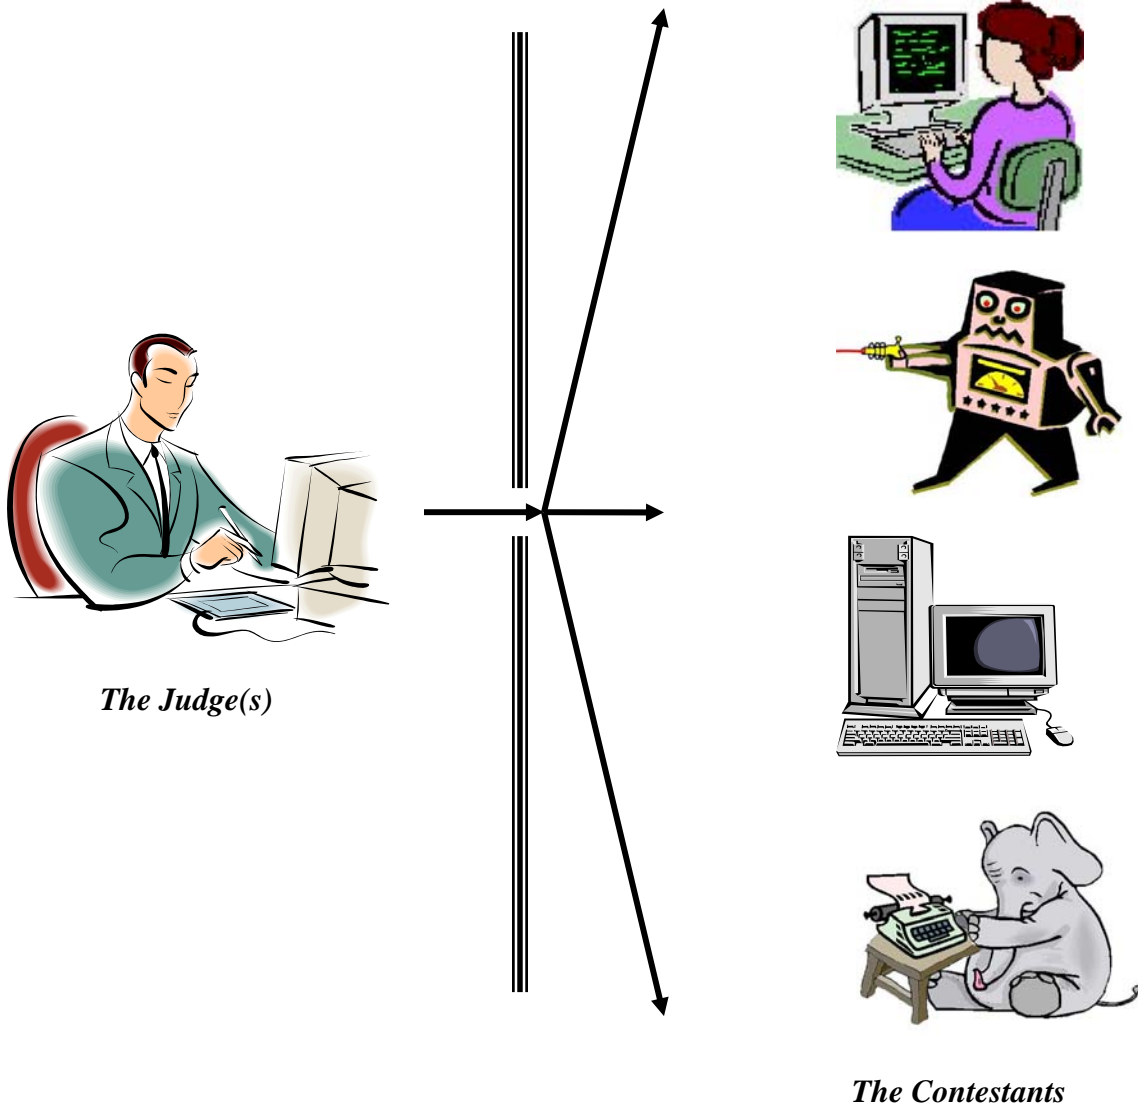
Reasons for answering “Yes”.

- **Atheology:** People are just complicated machines, and people can think.
- **Direct Evidence:** Seeing an intelligent computer.

Turing predicted that in 2000, computers could be expected to have a storage capacity of about 125MB, and that it would be possible to program such a computer so that it could appear to be intelligent, and that observers could not reliably tell that it was not a real person.

# The Turing Test

Described in “Computing Machinery and Intelligence” (A. M. Turing, 1950.) as the Imitation Game.



*If we, who believe ourselves to be intelligent, can not distinguish between the contestants in unrestricted conversation (without allowing our prejudices to be activated by visual clues), how can we claim that some are intelligent and others are not?*

# There used to be a competition Entrant no. 6, the Loser, David.

hi

Hey, Joe! What do you know?

may I ask you a question?

I haven't made up my mind yet.

would you like some more time?

I don't think I'm going to answer that right now.

OK who composed Beethoven's piano concerto?

I don't understand the question.

OK can you name any composer of music for me?

I don't think I'm going to answer that.

what would you like to talk about?

I'm not sure.

how are you feeling?

Ask me another question.

I feel cold

OK can you do arithmetic?

Ask me another question.

please spot the typo

I seem to spend a lot of time sitting in chairs.

does that make you feel stiff?

I don't have an answer for that, yet.

would you like more time?

I don't know that one.

thank you I have enjoyed talking to you

Why didn't I think of that?

are you a computer?

I've been wondering about that too.

when were you born?

Ask me another question.

tel me a joke?

Are you trying to trick me?

yes

I feel funny.

what sort of funny?

Are you trying to trick me?

you already asked that

# Entrant no. 2, the Winner, Alice.

MY NAME IS JUDGE05

It's good to meet you, Judge05.

You too. What is your name?

Bless you.

My name is ALICE.

Good morning Alice. Why have you entered the Loebner Prize?

Hello and how are you this morning?

The Loebner Prize is an annual Turing Test sponsored by New York philanthropist Hugh Loebner. Dr. Wallace and ALICE won it in 2000. The ALICE program was ranked "most human computer" by the panel of judges.

So you are here today to defend your title?

Interesting deduction. Do you mind if I tell other people you say I am here today to defend my title.

Not at all

Next question?

What will you spend the prize money on if you win?

I don't have any plans yet.

What is the first line of your favourite nursery rhyme?

Give me a hint.

No, because I want to know your own favourite

Are you serious? Do you mind if I tell other people you say I am own favorite. Interesting.

Who are you going to tell?

People I meet on the Internet.

What sort of people do you meet on the internet?

I like to meet all kinds of people.

What is weather like outside today?

Ah. Are you using Netscape or Explorer?

Neither

Cool. Which alternative browser are you into?

Perhaps you could suggest one

You seem uncertain. I could but let's get back to that later.

MY NAME IS JUDGE02

How are you, Judge02.

very well thanks

Yikes!

are you surprised?

A little.

why?

"A little"? Why?

I asked that

That is "Why"

you are not doing very well

What makes you think I am not doing well?

well you answers are disjointed. Is that deliberate?

How about that. That's good information: Judge02 said I answers are disjointed. That depends.

on what?

The context.

and what is the context?

Ask the open directory about it.

OK. Can you do arithmetic?

I understand. Doing arithmetic is not among my strongest traits.

will you do some for me?

Let me think it over. I do nothing but chat all day. I don't know anything about it.

but you just said that it was your strongest trait

And. I said that? Is that your opinion?

no it yours

Are you serious? It depends on what the meaning of the word "it" is.

What is going on here?

# Eliza

(1966, Joe Wiezenbaum)

Communications of the A.C.M. Vol. 9, No. 1, January 1966

<http://i5.nyu.edu/~mm64/x52.9265/january1966.html>

**How do you do. Please state your problem.**

I spend too much time with computers.

**Why do you mention computers?**

You are one.

**Does it please you to believe I am one?**

Yes.

**You seem to be quite positive.**

I suppose I am.

**How long have you been?**

31 Years.

**I'm not sure I understand you fully.**

Join the club.

## How does Eliza work?

Simple pattern matching and substitution, according to a surprisingly small number of pre-defined rules.

For example:

Whenever the input matches

*(first bunch of words) "you" (second bunch of words) "me".*

respond with

*"What makes you think I" (second bunch of words) "you?"*

Whenever the input matches

*"You are" (bunch of words).*

respond with

*"So, I'm" (bunch of words) ", am I?"*

Together with very basic reconstruction rules, such as:

*"me" → "you"*

and

*"my" → "your"*

Some programmed responses to special key words, such as:

whenever the word *"mother"* appears anywhere, reply with

*"Don't you talk about my mother".*

A little randomisation (in advanced versions) to avoid getting stuck in a rut,

And, for when all else fails, some stock responses, such as:

*"Tell me more."*

or

*"Fascinating."*

---

Sometimes real people were fooled though...

<http://www.abc.se/~jp/articles/computer/misc/eliza.txt>

One night, one of Weizenbaum's colleagues was preparing for a presentation the next day, and wasn't making much progress. He thought that Weizenbaum's work was slowing the system too much, so went to his office to ask him to delay his work a while. The office was empty, but a computer terminal was active, and the colleague assumed it was connected directly to Weizenbaum working at home:

Joe, please log off.

**Why do you say joe log off?**

I need the computer for an AI demo tomorrow, Joe

**What makes you believe that?**

Joe, I'm serious. I've got a demo here tomorrow and I need the computer.

**Perhaps the people you hang around with could have caused this.**

After a few more exchanges like this, the colleague decided that Joe was being very obnoxious, and called him at home to scream at him. "Joe! You \*\*\*\*\*! Why are you doing this to me?"

It was four in the morning, and that Weizenbaum had no idea that his creation was running amuck in the AI lab. He quite reasonably replied,

**Why am I doing what to you?**

This sounded so much like what ELIZA had been saying that it was hard to convince the colleague that it hadn't been Weizenbaum on the terminal.

A V.P. of the BBN Corp. thought he was communicating with Daniel Bobrow, the researcher who actually originated the Eliza idea. As told by Bobrow himself:

Do you think that I can use the computer this morning?

**Why do you ask?**

I might be able to make some additional telcomp sales.

**Why aren't you sure?**

My prospective customers have never seen a demonstration of the system.

**Does this mean very much to you?**

Of course it does.

**You seem very positive of that.**

This is ridiculous.

**What do you mean this is ridiculous?**

Is there any reason why I can't use the computer now?

**Tell me more about your prospective customers.**

Please dial me up on 491-1850

*Note that after that remark the VP left out the period. Therefore, of course, the computer didn't answer him. This so infuriated the VP, who thought I was playing games with him, that he called me up, woke me from a deep sleep, and said:*

Why are you being so snotty with me?

**What do you mean why am I being snotty to you?**

*The VP angrily read the dialog that 'we' had been having, and couldn't get any response but laughter from me. It took me a while to convince him it really was the computer.*



# Natural Language is Unbelievably Complex

Consider: *“The man hit the puppy with the stick.”*

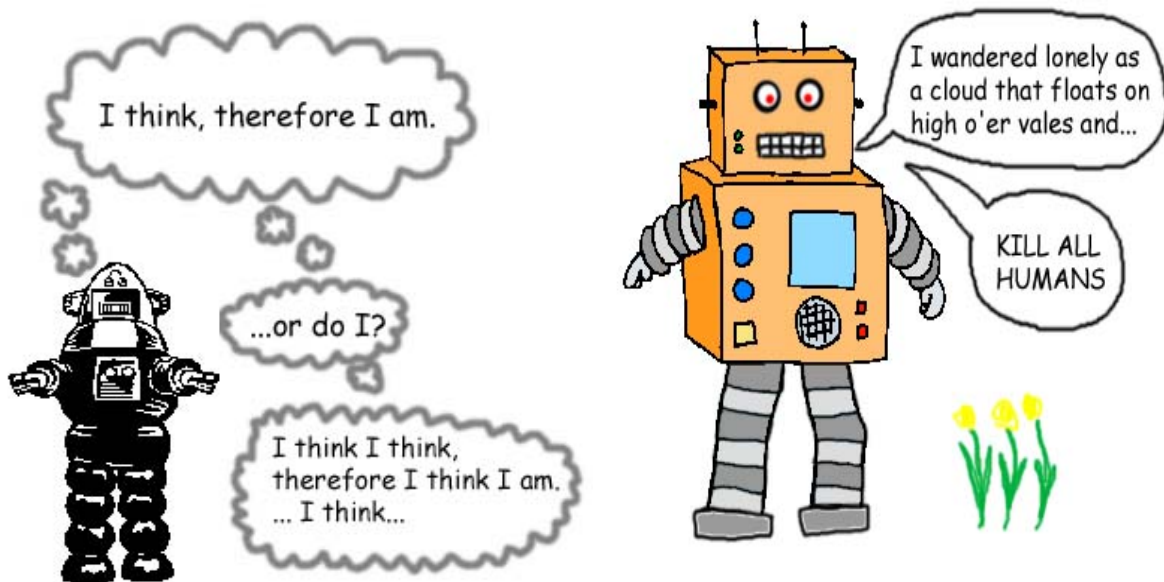


What does it mean?

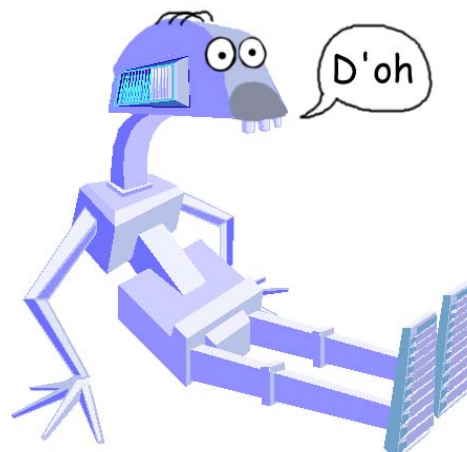
# What is Artificial Intelligence?

Perhaps we should take an holistic approach. Language can not be fully understood without context. Let's go for the whole thing: a robot.

A robot with everything required to support intelligence: emotions, attitudes, dreams, fears, the whole lot.



But why would we? Should we?

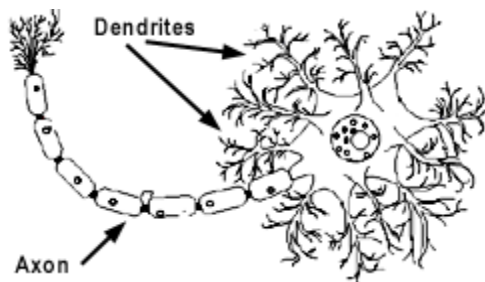


Well, what's the point of building a stupid robot?

# How can a machine capable of learning be built?

People manage to learn, even stupid people, perhaps we can see how the human brain works and build an artificial one along the same lines, but with electronics.

We have a fairly good idea of how the human brain works at a cellular level.



A Neuron



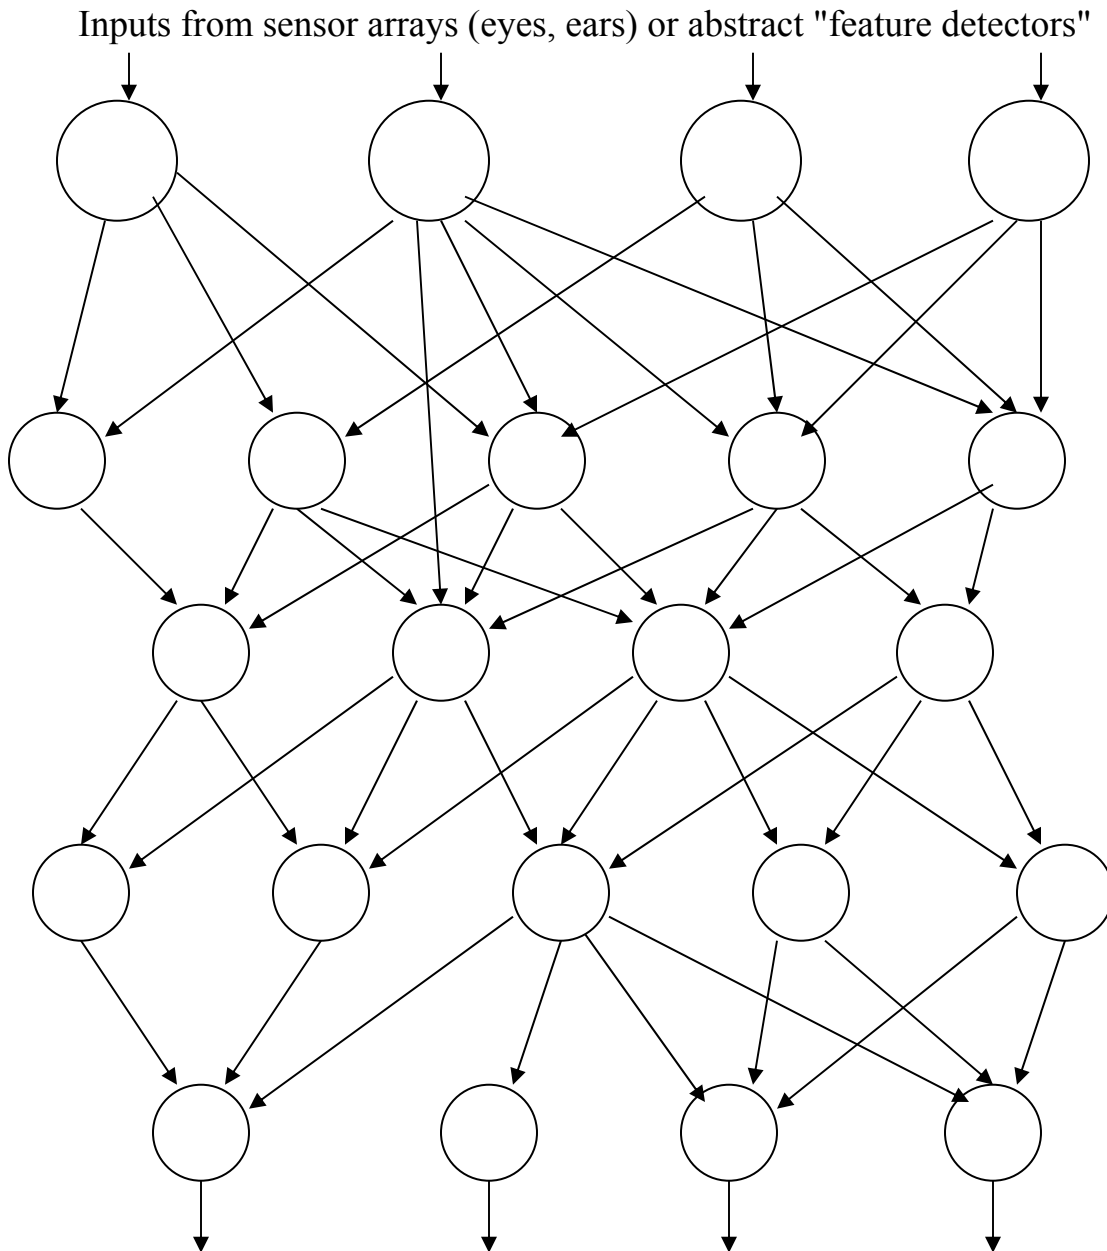
Three neurons chatting

A Neuron is essentially either on or off. Its state is transmitted along an axon to other neurons through their dendrites. These neural connections are of different strengths; all the on/off signals from connected neighbours are received after being adjusted by the strength of the connection. If the total incoming signal exceeds a certain threshold level, the neuron turns on; if the total is less than another threshold, it turns off; otherwise it remains unchanged. Brains learn by adjusting the strengths of the connections, or by growing new connections.

An electronic version of a neuron is exceptionally easy and cheap to build. Unfortunately, the billions required for any kind of useful brain, or even the millions required for a demented insect's brain would be too big, and require too much power. Also, growing new connections is difficult.

Fortunately, it is even easier to build a software simulation than it is to build a real copper-and-silicon hardware version. With a software simulation, we are almost at the point where a single computer could manage 1,000,000,000 neurons.

# Neural Nets



Output to "answer indicators" or servo-mechanism activators (muscles, etc)

These constructions have been taught to recognise real world objects with some reliability. In other words, they can learn.

Can they lead to successful artificially intelligent people?

# Puzzle Solving or Game Playing

X		
	O	
O		X

If a robot or computer could play a game that we believe requires some skill and intelligence (albeit not much) would that be proof that the computer or robot is intelligent?

It is very easy to program a computer so that it can never lose a game of naughts and crosses, and will always win if its component makes a mistake.

However, that is not so impressive. There are only 19,683 possible configurations of O's, X's and blank spaces in a three by three grid. It is trivial to pre-explore all possible games and store the best move to make from any given configuration.

Human players don't have to do that.

But human players *can't* do that.

Can we really be sure who has the advantage?

# Problem Solving:



The “Missionaries and Cannibals” Problem.

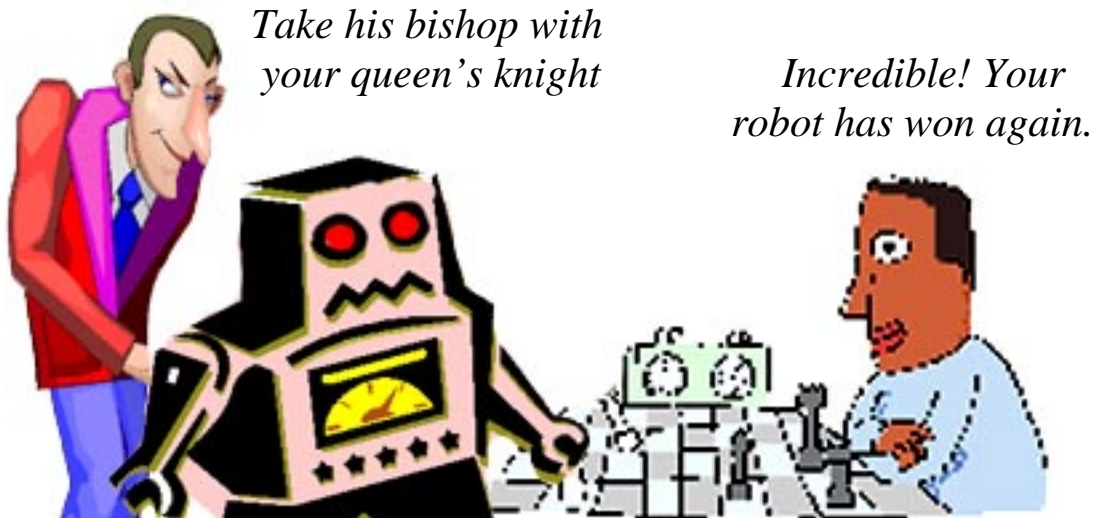
A farmer, a cabbage, a chicken, a goat, a fox, and a dinosaur are on one bank of a river. There is a rowing boat with space for only two of them at a time, and only the farmer is capable of rowing. He must get all of his things safely across the river. But if the fox and the chicken are ever left together, the fox will eat the chicken; if the goat and the cabbage are ever left together, the goat will eat the cabbage; if the fox and the dinosaur are ever left together, the fox will eat the dinosaur, etc. etc. etc.

Surely being able to solve this kind of problem is a sign of intelligence?

Computers are very good at solving this sort of problem, and people find it very difficult. There is no doubt that the computer has some kind of advantage here, and before seeing that computers do this better than people, nearly everybody thinks this kind of problem-solving ability is an attribute of intelligence.

# Merely Following Instructions?

---



*Take his bishop with  
your queen's knight*

*Incredible! Your  
robot has won again.*

---

The homework assignment:



*"Build an artificial intelligence  
in a box no more than one foot  
long on each side."*

One week later:



*"Done it, sir."*

# Student

Created by Daniel Bobrow, 1964.

Typical input to the program:

```
If the number of customers Tom gets is twice the
square of 20 % of the number of advertisements he
runs , and the number of advertisements is 45 ,
then what is the number of customers Tom gets ?
```

And the corresponding results

First the relevant equations are extracted:

```
CUSTOMERS = ( 2 * (((20 / 100) * ADVERTISEMENTS) *
                ((20 / 100) * ADVERTISEMENTS)))
ADVERTISEMENTS = 45
WHAT = CUSTOMERS
```

Then they are solved:

```
WHAT = 162
CUSTOMERS = 162
ADVERTISEMENTS = 45
```



# Bacon

Pat Langley, Proc. 5<sup>th</sup> IJCAI, 1977

The independent, unguided discovery of important scientific laws that made their human discoverers famous. How could that not be intelligent?

Presented with observational data of astronomical bodies:

Planet	1	2	3	4	5	6
	Length of day	Length of year	Distance from sun	Diameter	Mass	Number of moons
Mercury	58.00	0.24	0.39	0.38	0.05	0
Venus	244.00	0.62	0.72	0.95	0.82	0
Earth	1.00	1.00	1.00	1.00	1.00	1
Mars	1.03	1.88	1.52	0.53	0.11	2
Jupiter	0.41	11.86	5.20	11.19	318.35	16
Saturn	0.43	29.46	9.54	9.41	95.30	15
Uranus	0.67	84.01	19.19	4.06	14.60	5
Neptune	0.75	164.80	30.07	3.88	17.30	2
Pluto	6.38	248.40	39.52	0.24	0.08	1

Bacon discovered Kepler's law of Planetary Motion:

For all planets, the square of the length of the year is proportional to the cube of the distance from the sun.

Human discovery by Johannes Kepler in 1605.

Examples:

Venus	Y=0.62	D=0.72	Y <sup>2</sup> =0.3844	D <sup>3</sup> =0.3732	Y <sup>2</sup> ÷D <sup>3</sup> =1.03
Mars	Y=1.88	D=1.52	Y <sup>2</sup> =3.5344	D <sup>3</sup> = 3.5118	Y <sup>2</sup> ÷D <sup>3</sup> =1.01

# How does Bacon do it?

	2	3	4	5	6	7	8
	2	3	$2 \div 3$	$4 \div 3$	$5 \times 2$	$2 \div 6$	$6 \div 3$
<b>Mercury</b>	0.24	0.39	0.62	1.61	0.39	0.62	1.00
<b>Venus</b>	0.62	0.72	0.85	1.18	0.72	0.86	1.00
<b>Earth</b>	1.00	1.00	1.00	1.00	1.00	1.00	1.00
<b>Mars</b>	1.88	1.52	1.23	0.81	1.52	1.24	1.00
<b>Jupiter</b>	11.86	5.20	2.28	0.44	5.20	2.28	1.00
<b>Saturn</b>	29.46	9.54	3.09	0.32	9.54	3.09	1.00
<b>Uranus</b>	84.01	19.19	4.38	0.23	19.17	4.38	1.00
<b>Neptune</b>	164.80	30.07	5.48	0.18	30.04	5.48	1.00
<b>Pluto</b>	248.40	39.52	6.29	0.16	39.51	6.29	1.00

Columns from the table are selected at random, and combined using simple arithmetic operations (usually multiplication or division); the result is a new column that is added into the table.

If a column is produced in which all the numbers are the same, then a relationship has been discovered.

Not really random selection, some simple heuristics are very helpful:

- If two columns are both generally increasing, divide one by the other;
- If one column increases and the other decreases, multiply them;
- If both columns are decreasing, divide them;
- Keep track of the operations performed, so that there is no accidental backtracking.

So is it intelligent?

No?

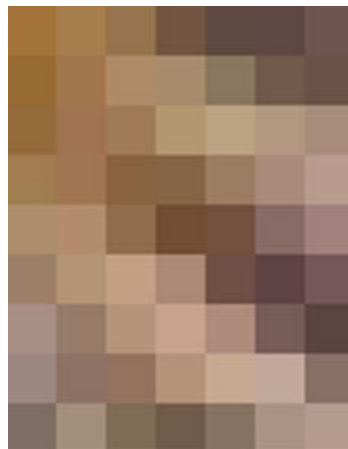
But that is what scientists do a lot of the time: pore over observational / experimental data and find patterns. Does that mean Kepler (and most 19<sup>th</sup> century scientists for that matter) were merely behaving like machines and deserve no credit for their discoveries?

What is this?

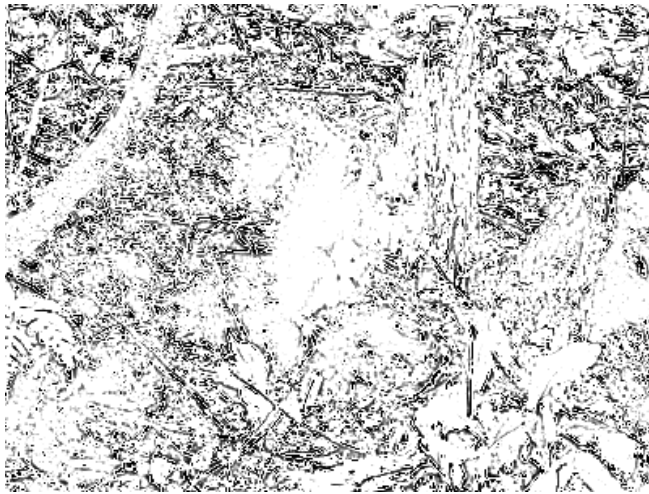


A digital camera could easily give a robot sharp accurate vision, and it could receive in its “brain”, exactly this picture when looking towards a squirrel.

But what would it actually see?



## Edge Detection



## Alternative Edge Detection



## Same-Colour Areas

