Can a learning robot solve a 2-D jeep problem?

Akira Imada

Brest State Technical University

Belarus

Looking for "a-needle-in-a-haystack"

A needle hidden in a huge grid world

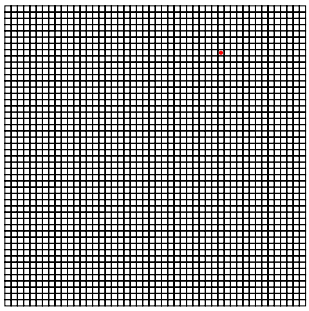

Can a navigation robot reach it?

$\Downarrow$

A 2-D version of a-needle-in-a-haystack.

What is a-needle-in-a-haystack

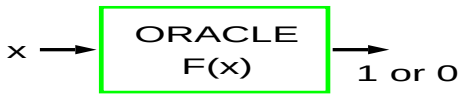

"Find x (among N items) such that $F(x) = 1$"

when

only x fulfills $F(x) = 1$ while $F(x) = 0$ for all others.

(We might call $F(x)$ a fitness function of this problem.)

$O(N)$ steps are necessary!

How we search the needle as efficiently as possible?

$\Downarrow$

More efficient than a Random/Exhaustive search?

Hinton & Nowlan's Needle (1987)

- A-needle $\Rightarrow$ Just one configuration of 20-bit binary strings.
- Haystack $\Rightarrow$ $2^{20} - 1$ search points.

Their choice of 20-bit was a good one!

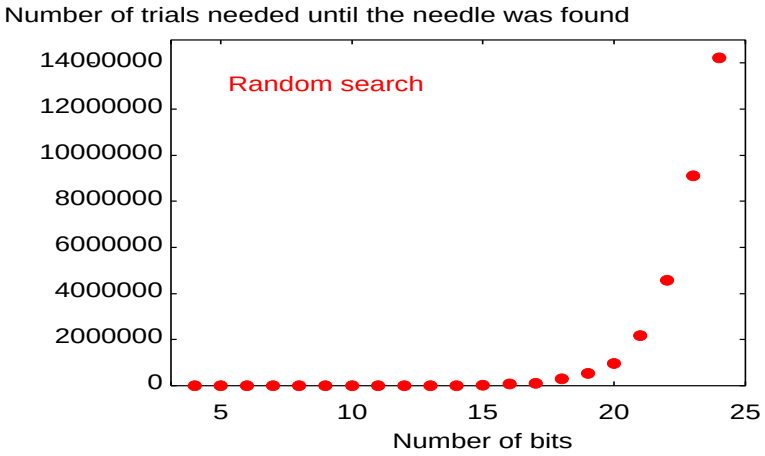

Hence, their experiment was not a very difficult one.

How we evolve all-fitness-0 chromosomes?

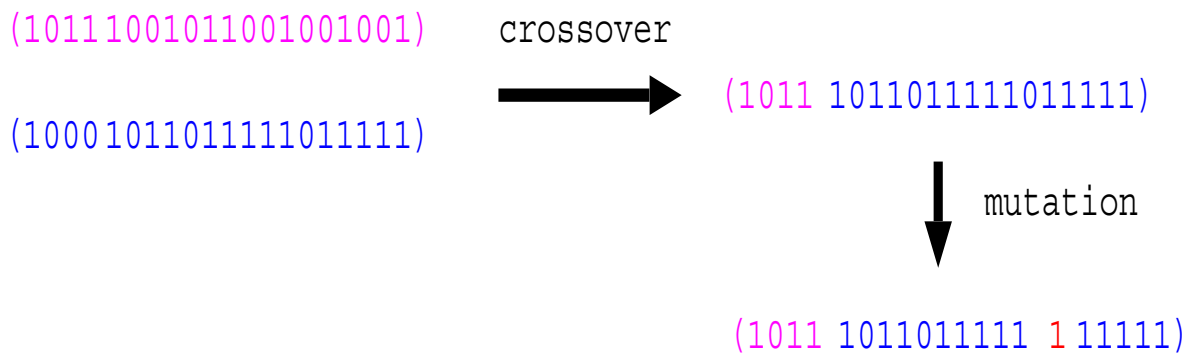

- What about the task called "All-one-problem?"

  $(9 \times 14 \Rightarrow 16 \Rightarrow 17$ vs. $0 \times 0 \Rightarrow 0 \Rightarrow 0)$

Hinton & Nowlan's brilliant trick

Evolution under Baldwin effect

— Lifetime learning of phenotype

A genotype:

(1090109901100100109 1)

Its phenotype:

(1010100101100100100 1)
(1000101101100100101 1)
...

Their assumption

The closer the *genotype* to the needle,

the faster the learning of *phenotype.*

$\Downarrow$

This makes the sharp-fitness-peak a more smooth one.

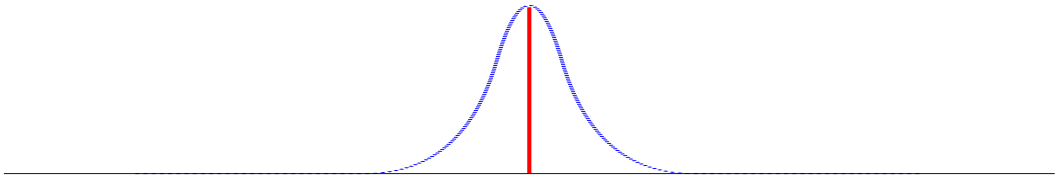

Baldwin effect looks great!

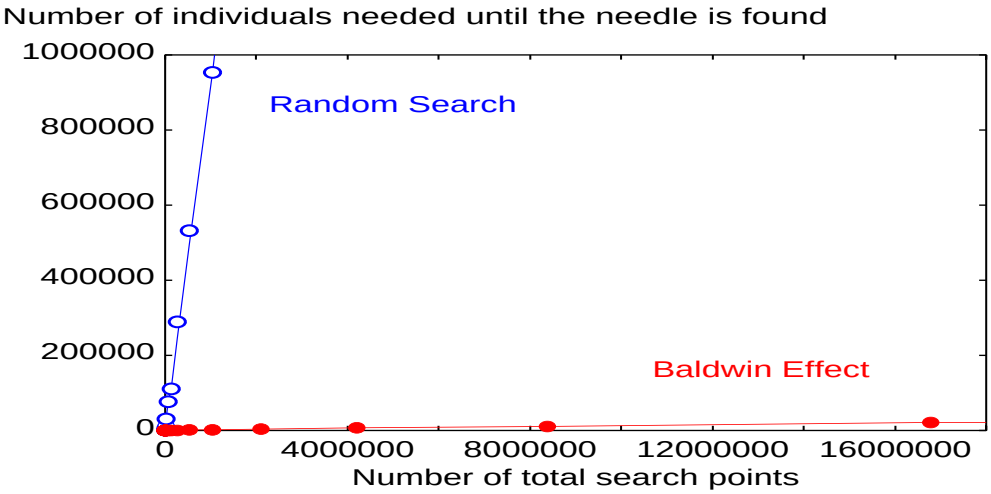

Are we really happy with this?

Why should we continue evolution
when lifetime-learning already has found the needle?

⇓

(Turney 1987)

"Not from an engineering but a biological interest."

Still we are not so happy.

The number of individuals, or *the number of total points visited*?

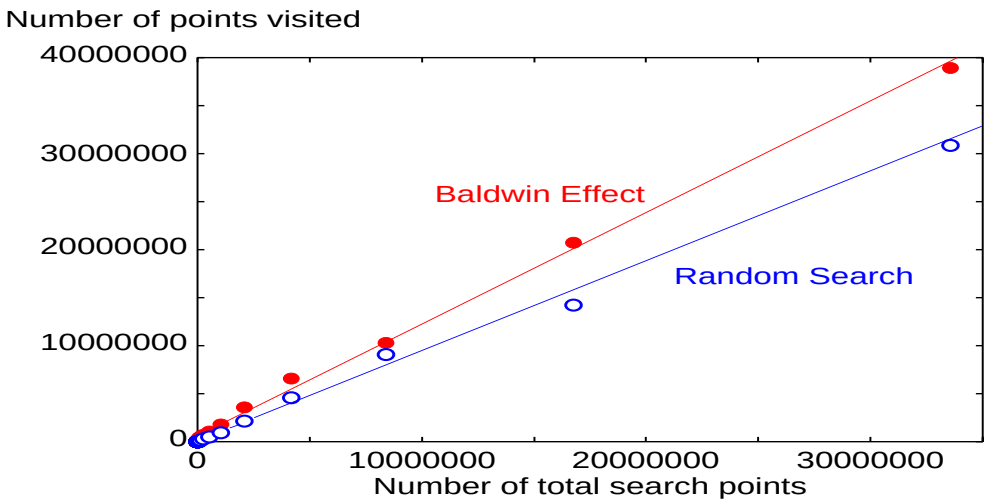

Effect of like-to-hear-what-we'd-like-to-hear?

An intelligent strategy is sometimes worse than a random one!

A random walk to the needle

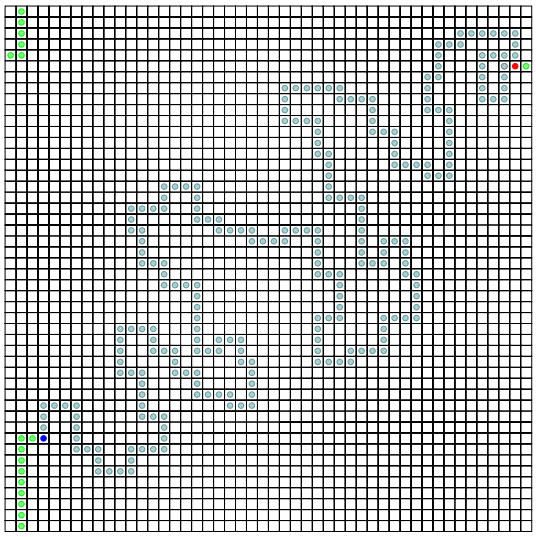

Random walk evolved to be minimized

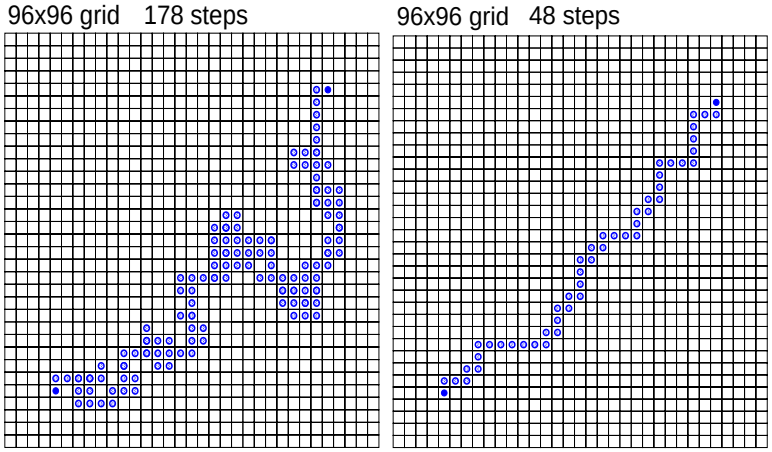

Evolution successful

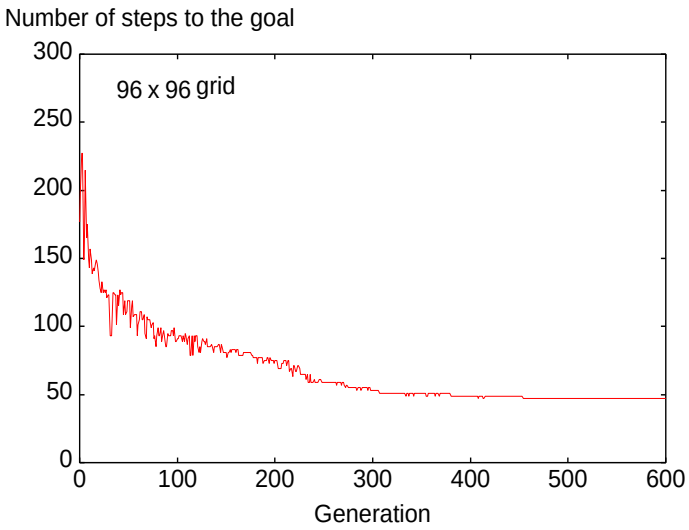

Number of steps grows linearly

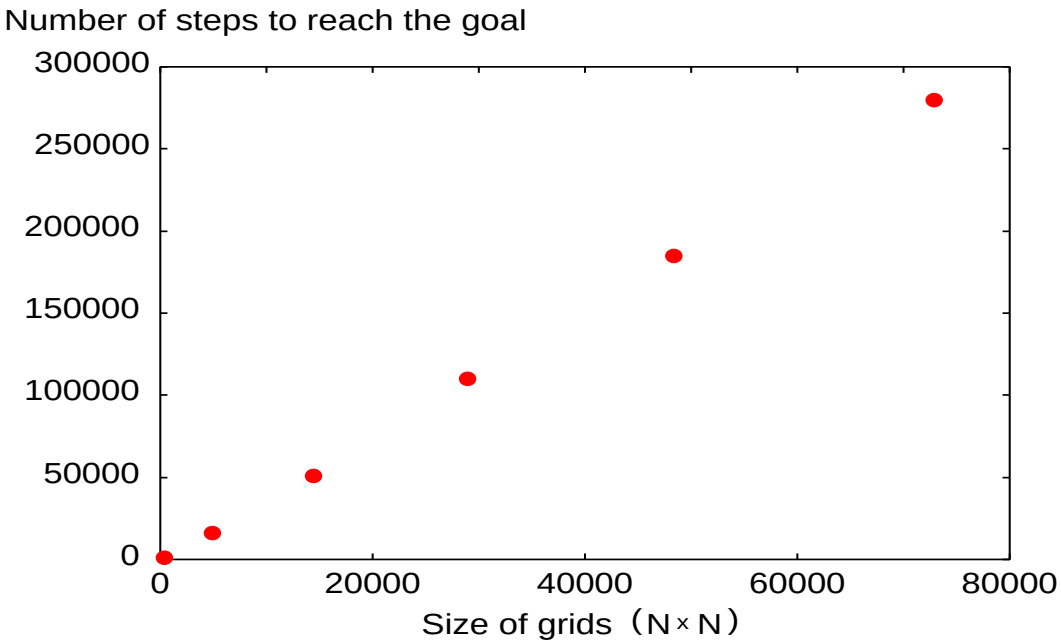

What if the grid is huge?

It seems less likely to reach the needle by chance.

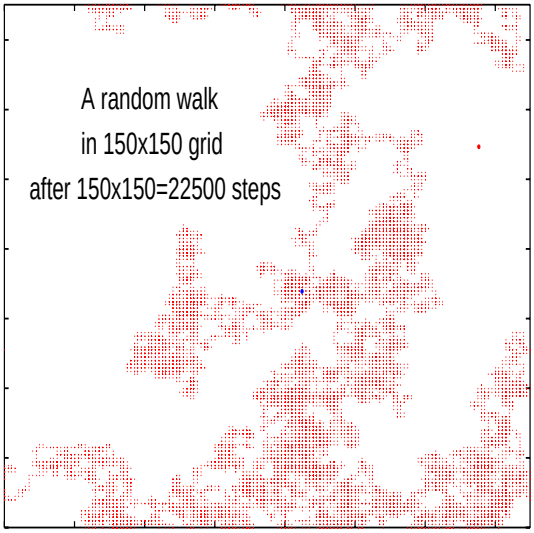

Evolution will be harder as grid becomes larger

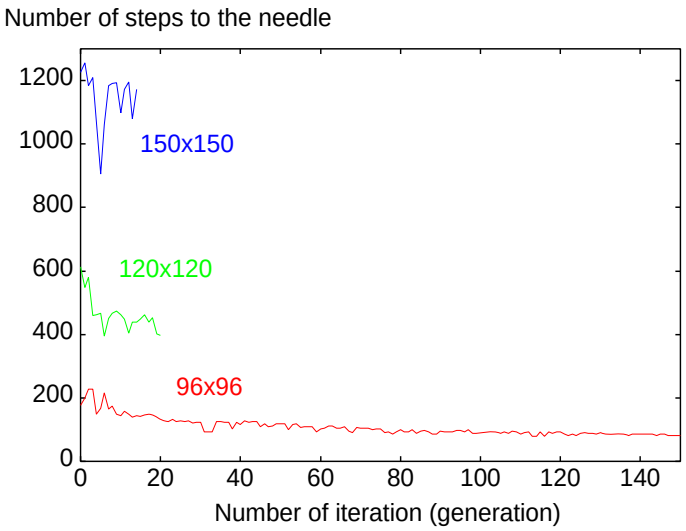

A success at a small scale
is not a royal road to a real success.

The Jeep Problem

The 52nd problem

in the

"Propositiones ad acuendos inventes" (in Latin)

attributed to

Alcuin of York (732–804)

$\Downarrow$

Camel carrying grain in a desert.

The task in the modern version

Maximize the distance a Jeep can penetrate into 1-D desert.

- Jeep can move a unit distance with a unit fuel.
- Jeep can unload its fuels anywhere in the desert.
- Fuels can be filled only at the base.
- Jeep can go back to the base n times to re-fill its tank.

An example of a success

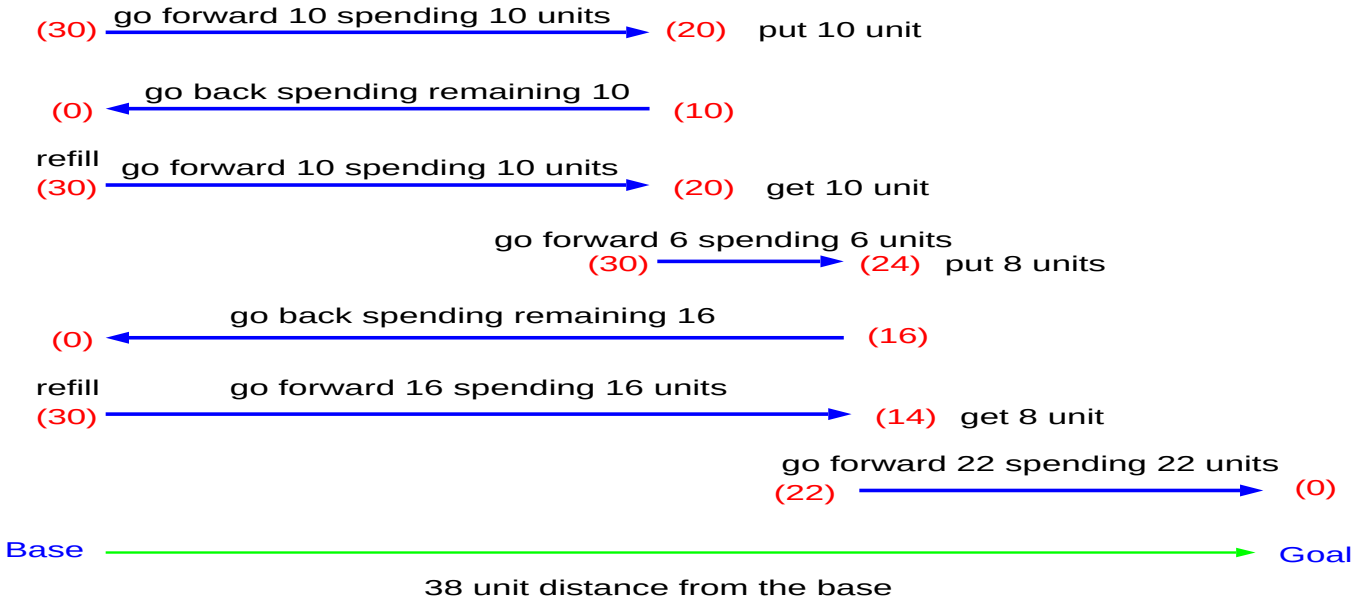

An analytical solution

by N. J. Fine (1947)

$$d_0 = 0, \quad d_n = d_{n-1} + 1/(2n-1).$$

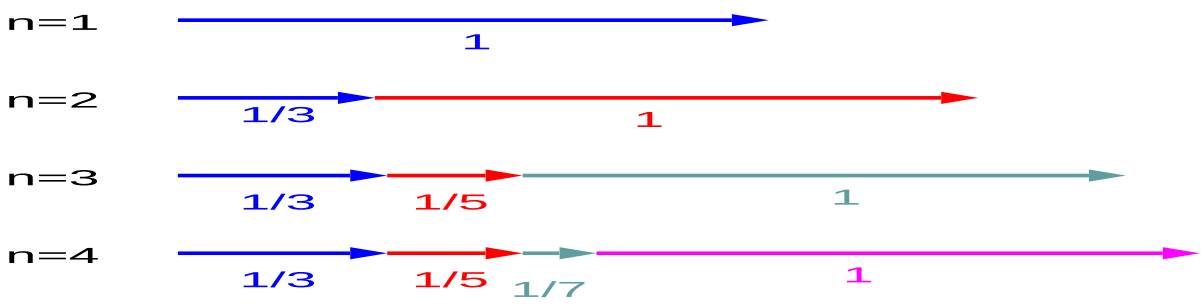

The optimum solution for n=3

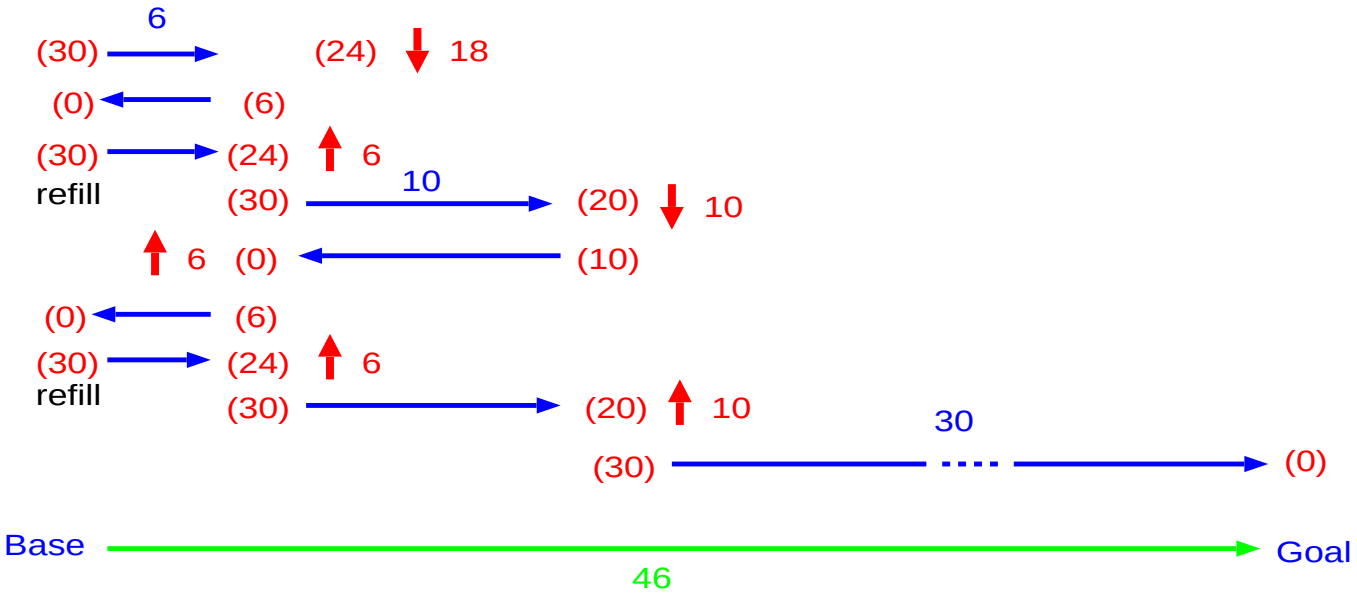

A solution by a GA

Przemysław Klęsk (2004)

A feasible candidate

Not feasible

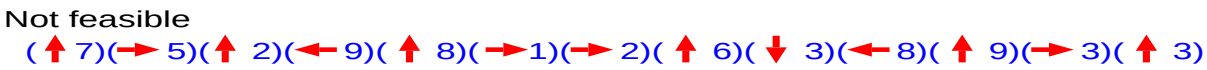

Optimal solution, e.g., for n=2

Conluded, "Solution for $n > 3$ was not available!!

An extension to 2-D Desert

"Maximize the penetration" or "Find the needle"
by repeating the procedure:

(1) Start the base.

(2) Navigate the desert.

(3) Put fuels somewhere, or find the fuels to get.

(4) Return to the base.

Examples of a possible trace

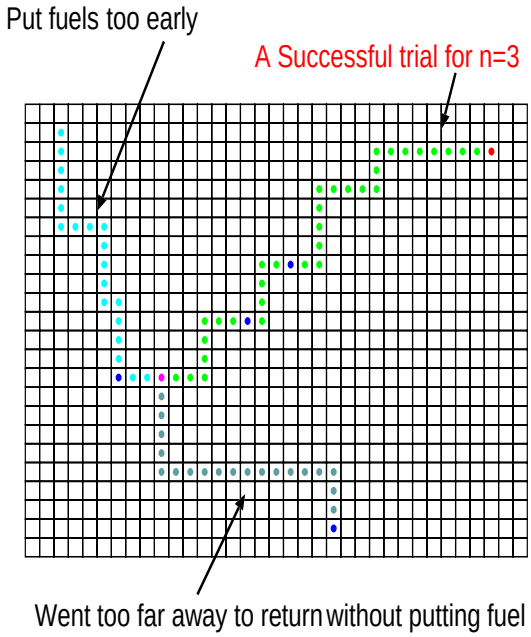

Still resists to be solved!

Evolutionary Computation

.....

Reinforcement learning

etc

$\Downarrow$

An extremely tough benchmark

Grover's quantum search (1997)

A speed up from $O(N)$ to $O(\sqrt{N})$.

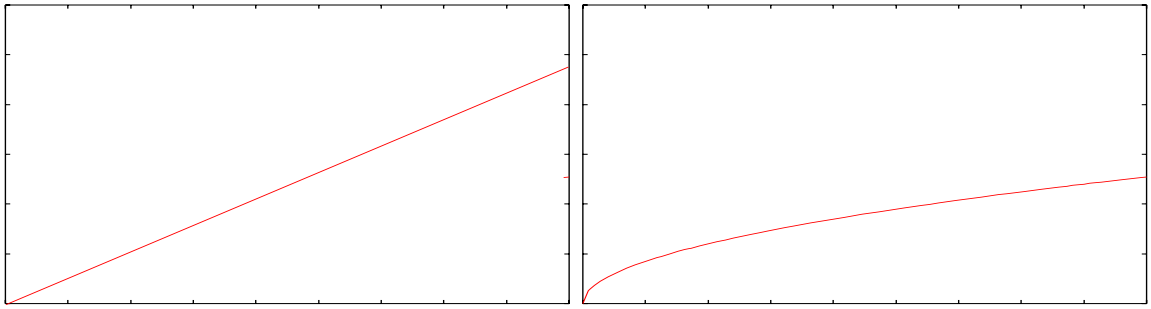

A strange path of quantum computation

When a particle goes from A to B,

it takes all possible paths at the same time.

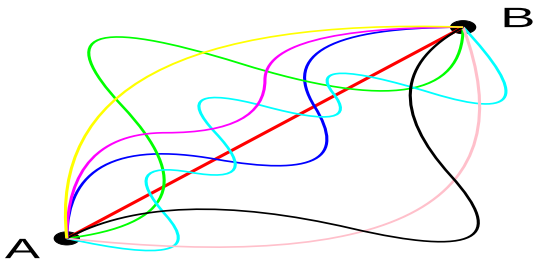

Quantum Robot

Paul Benioff (2002)

$\Downarrow$

"Only faster than classical search in more than 2-D!"

Summary

- Two challenges:
  - ⋆ A needle hidden in a huge grid world.
  - ⋆ 2-D Jeep Problem by a learning navigation robot.

    Can a robot learn to solve this? Then how, if any?
- We have to be careful about an effect of
  tend-to-hear-what-we-want-to-hear.
- Success at a small scale is not a success in a real world.