How can we design a more intelligent path-planning?

Akira Imada

Brest State Technical University (Belarus)

Is AI really intelligent?

They might be more efficient and/or effective than human but still less intelligent, at least so far.

What is intelligence?

$\Downarrow$

Intelligence should be spontaneous, flexible, and/or unpredictable, more or less.

"I beg your pardon?"

Intelligent people try a different explanation for an easier understanding

while

others just repeat the same expression, maybe with a louder voice.

Let's make an agent to be spontaneous!

Behaviors might differ even in a similar situations.

Goal is

to make an agent behave differently even when it encounters the same situation as before.

Path Planning as a benchmark

An agent plans a shortest path from start to goal.

Genetic Algorithm (GA)

Agent decides actions following its chromosome

e.g., with chromosomes like
(4 1 1 2 1 3 2 3 3 4 2 1 3 4 3 2 3 3 2 1 1 3 2 1 3 ...)
indicates the agent a route to be followed.

Random walk evolved to be minimized

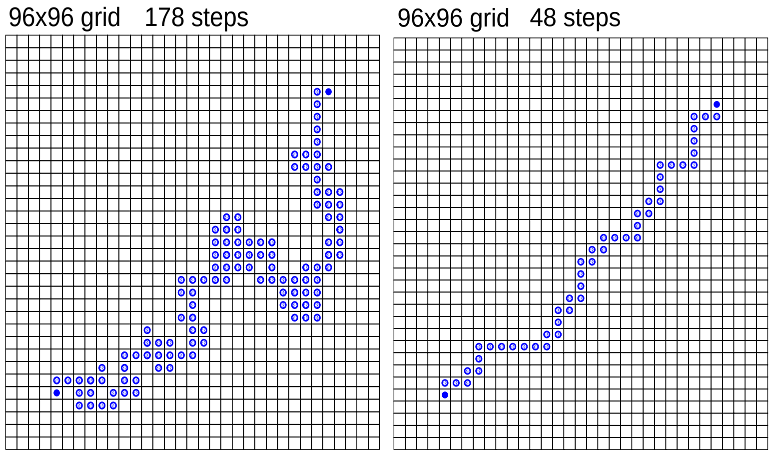

Task was easy

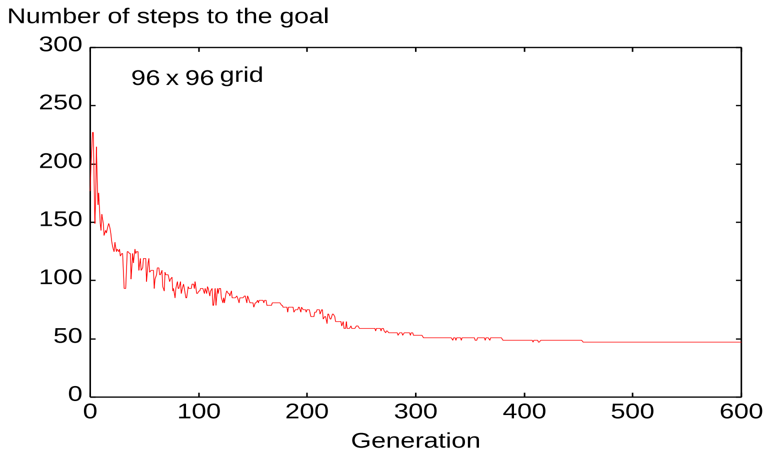

(An interpretation of GA from robotics point of view)

Edelman's Neural Darwinism (1978)

Assume evolution in an agent's brain in advance the action.
(e.g. 100 random chromosomes to be evolved)

$\Downarrow$

after a convergence

$\Downarrow$

Agent acts following the best chromosome.

Route can be different from run to run, if any.

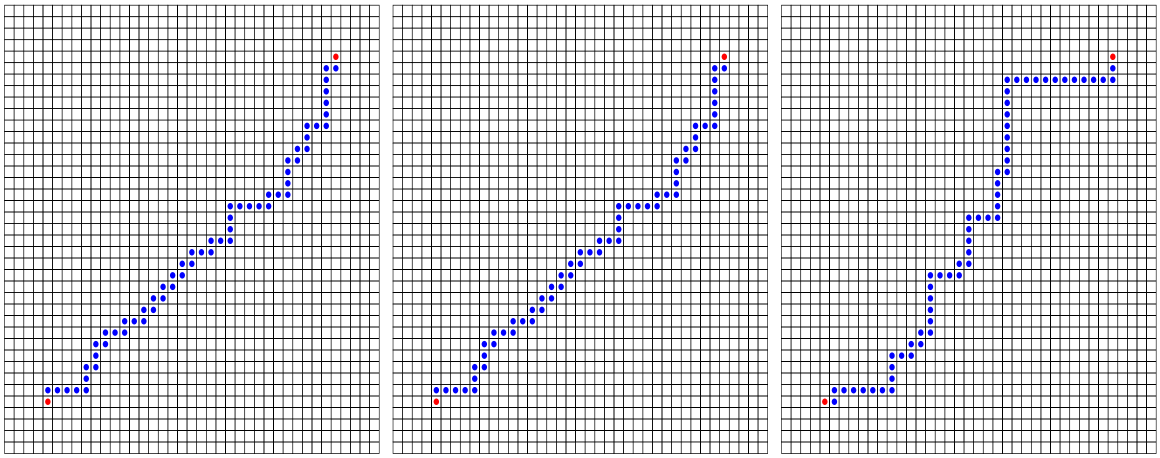

It's not so realistic, however.

No one, these days, believe this really happens.

It takes a too long time to be a reality if the task is hard.

$\Downarrow$

Let's dismiss GA now.

What about NN?

With fixed weights
$\Downarrow$
An identical behavior in an identical situation.

What if we modify weights during action?

(Floreano et al. 2000)

(1)

Recurrent McCullough-Pitts Neurons with Sigmoid

Learning Rule

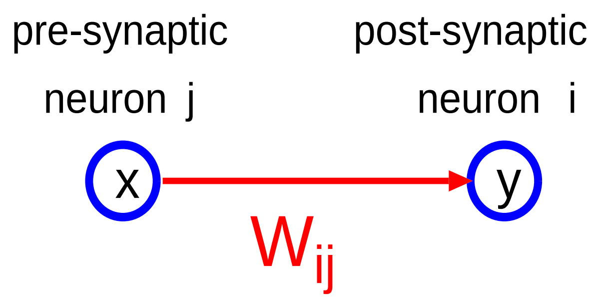

$$w_{ij}(t+1) = w_{ij}(t) + f(x_j(t), y_i(t)) \quad (x_j, y_i \in [0,1])$$

Optimum configuration of learning rules.

$$\eta_{ij} x_i y_j \quad \text{(Hebb, 1949)}$$

or

$$a_{ij} x_i y_j + b_{ij} x_i + c_{ij} y_j + d \quad \text{(Durr et al. 2008)}$$

Can agents take a different minimum route every time following how it feels?

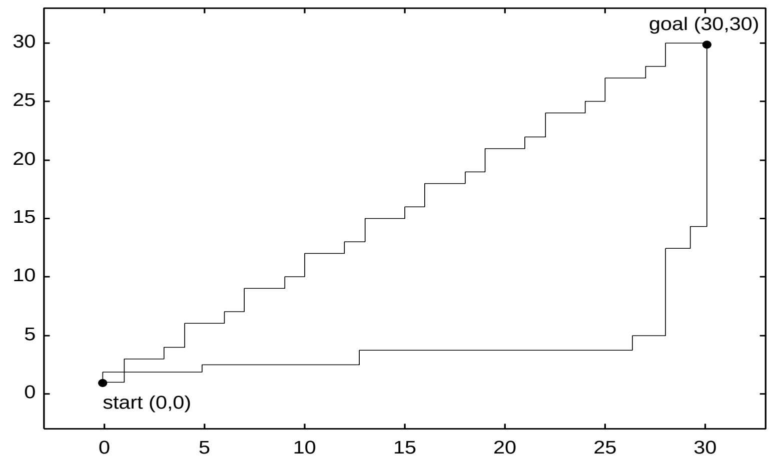

(2)
Recurrent Spiking Neurons

as a more biologically plausible model.

An exploration by Integrate & Fire Model

$$u_i(t) = u_r + (u_i(t - \delta t) - u_r)\exp(-\delta t/\tau) + \sum_j w_{ij} f_j(t - \delta t)$$

Can it learn so that path will finish at the goal with minimum length?

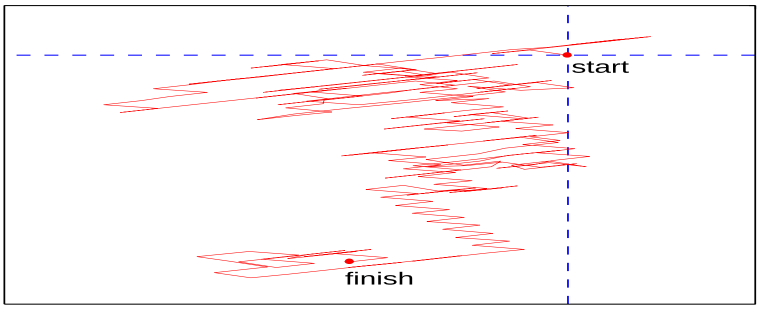

Spike timing dependent plasticity (STDP)

A counterpart of Hebb Rule of McCullough-Pitts Neurons

for Spiking Neurons.

What is STDP?

$$W(\Delta t) = \begin{cases} A_+ \exp(-\Delta t/\tau_+) & \text{if} \quad \Delta t \geq 0 \\ -A_- \exp(-\Delta t/\tau_-) & \text{if} \quad \Delta t < 0 \end{cases}$$

where $\Delta t = t_{post} - t_{pre}$

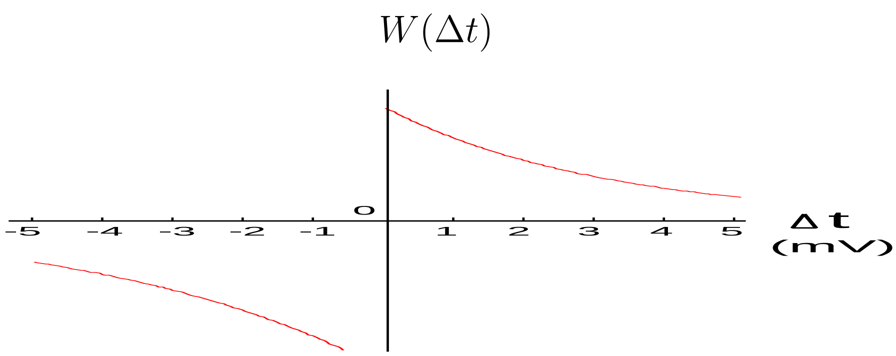

Reward-modulated STDP Learning
(Florian 2007)

$$w_{ij}(t+\delta t) = w_{ij}(t) + \gamma r(t+\delta t)\zeta_{ij}(t)$$

where

$$\zeta_{ij}(t) = P_{ij}^{+}(t)f_i(t) + P_{ij}^{-}(t)f_i(t)$$

$$P_{ij}^{+}(t) = P_{ij}^{+}(t-\delta t)\exp(-\delta t/\tau_{+}) + A_{+}f_j(t)$$

$$P_{ij}^{-}(t) = P_{ij}^{-}(t-\delta t)\exp(-\delta t/\tau_{-}) + A_{-}f_j(t)$$

Looks like a simple toy problem?

No!

More challenging benchmark

From one point to the same point maximizing its route with a limited fuel.

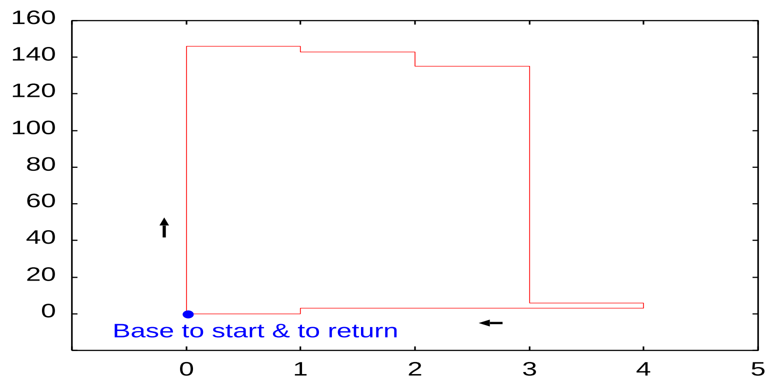

Brief Summary

We sometimes need spontaneous and unpredictable intelligence rather than efficiency or effectiveness like in case of SONY's AIBO

Robot dog vs. Real dog

When will a robot dog be more intelligent than real dogs?