How to evolve a neuron

W. Garrett Mitchener ◈ mitchenerg@cofc.edu
http://mitchenerg.people.cofc.edu

COLLEGE of CHARLESTON
DEPARTMENT OF MATHEMATICS
1770

Objective

An experiment in artificial biology: Using an abstract biochemistry model that is capable of general computation, simulate the evolution of a neuron. Agents should develop the ability to detect coincidences in an input spike train, and exhibit *spike-timing dependent plasticity (STDP).*

Virtual machine

Each agent consists of one cell. Its state is a vector `A[]`, where `A[t]=n` means there are `n` molecules of type `t` in the cell. Each cell's program consists of a list of instructions of the form

```
If A[pSwitch] ≥ threshold
then add 1 to A[pUp]
     and subtract 1 from A[pDown]
```

which form an *artificial regulatory network (ARN).*[3]

Genome

Each agent's diploid genome consists of two pairs of bitstrings (chromosomes). *Nucleotide bits* from the chromosome are translated using a majority-of-three code into *codon bits*, which are grouped and translated into integers using a Gray code.

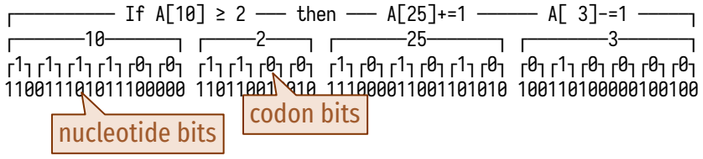

Coincidence detection task

Each agent is rated on its ability to process time series. During a time step when a *spike* is present in the input, 8 is added to `A[3]`. When `A[41]≥2`, the agent generates an output spike and 2 is subtracted from `A[41]`. When two input spikes arrive at almost the same time, the agent *should* generate an output spike. Otherwise, it *should not*. (Based on Mauthner neurons.[1])

Example coincidence detection task

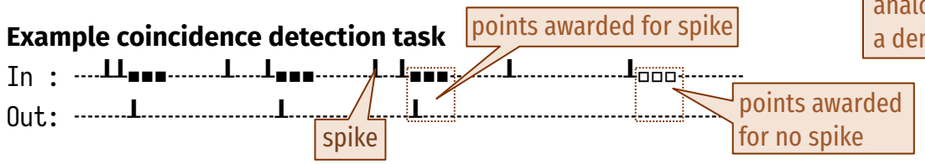

Hebbian learning: spike-timing dependent plasticity

After a period of high activity, an agent should spike after two input spikes that are moderately separated in time. After a period of low activity, the agent should spike reluctantly, only after pairs of input spikes that are very close. (Based on the NMDA receptor.[2,4])

Example high activity task (strengthens synapse)

Example low activitiy task (weakens synapse)

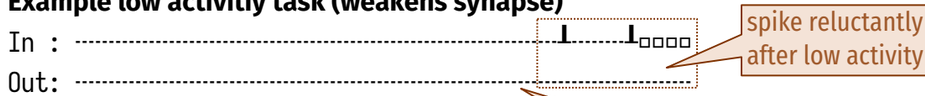

Total of 301 tasks. Length up to 527 time steps. Each generation takes ~15 seconds. Sample run of 20,000 generations takes 2 to 3 days.

Selection-mutation loop

The population consists of 1000 agents. Parents are picked using tournament selection. Their chromosome pairs experience crossover, gene deletion, gene duplication, and point mutations, resulting in haploid gametes. Two gametes combine to produce a diploid offspring. The lowest-rated 600 are replaced each generation. The simulation ends after the first perfectly scoring agent dies, or after 20,000 generations.

References

1. Korn & Faber, *Neuron* 41 (1) 2005
2. Li & Tsien, *New England Journal of Medicine* 361 (3) 2009
3. Mitchener, *Artificial Life* 20 (4) 2014
4. Sweatt, *Mechanisms of Memory, 2003*

Example perfect solution

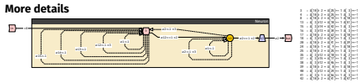

Simplified ARN diagram:

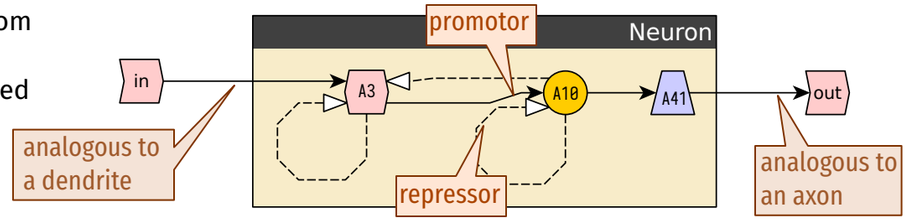

Varying the rating function

Some tasks seem more difficult, so should all be weighted the same? Let's vary between unscaled, weighted by increasing difficulty, and by decreasing difficulty. Perfect scores with any weighting are equivalent. 100 samples of each.