The Animat Path to Artificial General Intelligence

Claes Strannegård, Nils Svangård, David Lindström, Joscha Bach, Bas Steunebrink

AGA workshop IJCAI-17

Melbourne, 19 August, 2017

Our strategy for AGI

- Humans are animals. So human intelligence is a special case of animal intelligence.

- Eventually we want to reach human level AI, but first we have the challenge of X-level AI for other animals X.

- We are not done with X = insects. Going directly for X = human is like going for K2 when you are unable to climb a hill.

- Stewart Wilson
  - Coined the term animat = artificial animal
  - Wrote the paper *The Animat Path to AI*

- We follow this path aiming for general AI

Intelligent behavior

Intelligent behavior is to be repeatedly successful at satisfying one's [...] needs in diverse, observably different, situations on the basis of past experience

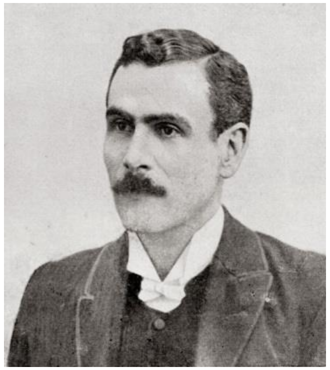

van Heerden, 1968

Applies to all animals with needs, not just humans!
Does not mention the notion of task.

Animal needs

- Water (osmoceptors)
- Energy (insulin receptors)
- Protein (amino acid receptors)
- Oxygen (CO_2 receptors)
- Integrity (nociceptors)
- Sleep (melatonin receptors)
- Proximity (pheromone receptors)
- Reproduction (sexual hormone receptors)
- Affiliation (oxytocin receptors)

Some animal needs and their associated interoceptors

Animal intelligence

- Nature is full of animals that are intelligent in van Heerden's sense, e.g. mammals, birds, reptiles, fish, crustaceans, and insects!

Mechanisms of animal intelligence

- Decision-making aimed at need satisfaction
- Reinforcement learning that strengthens/weakens behavior associated with reward/punishment
- Hebbian learning: ”if they fire together, they wire together”
- Sequence learning, which is a form of Hebbian learning with one of the signals delayed
- Forgetting: ”use them or lose them”

Which animal species use these mechanisms?

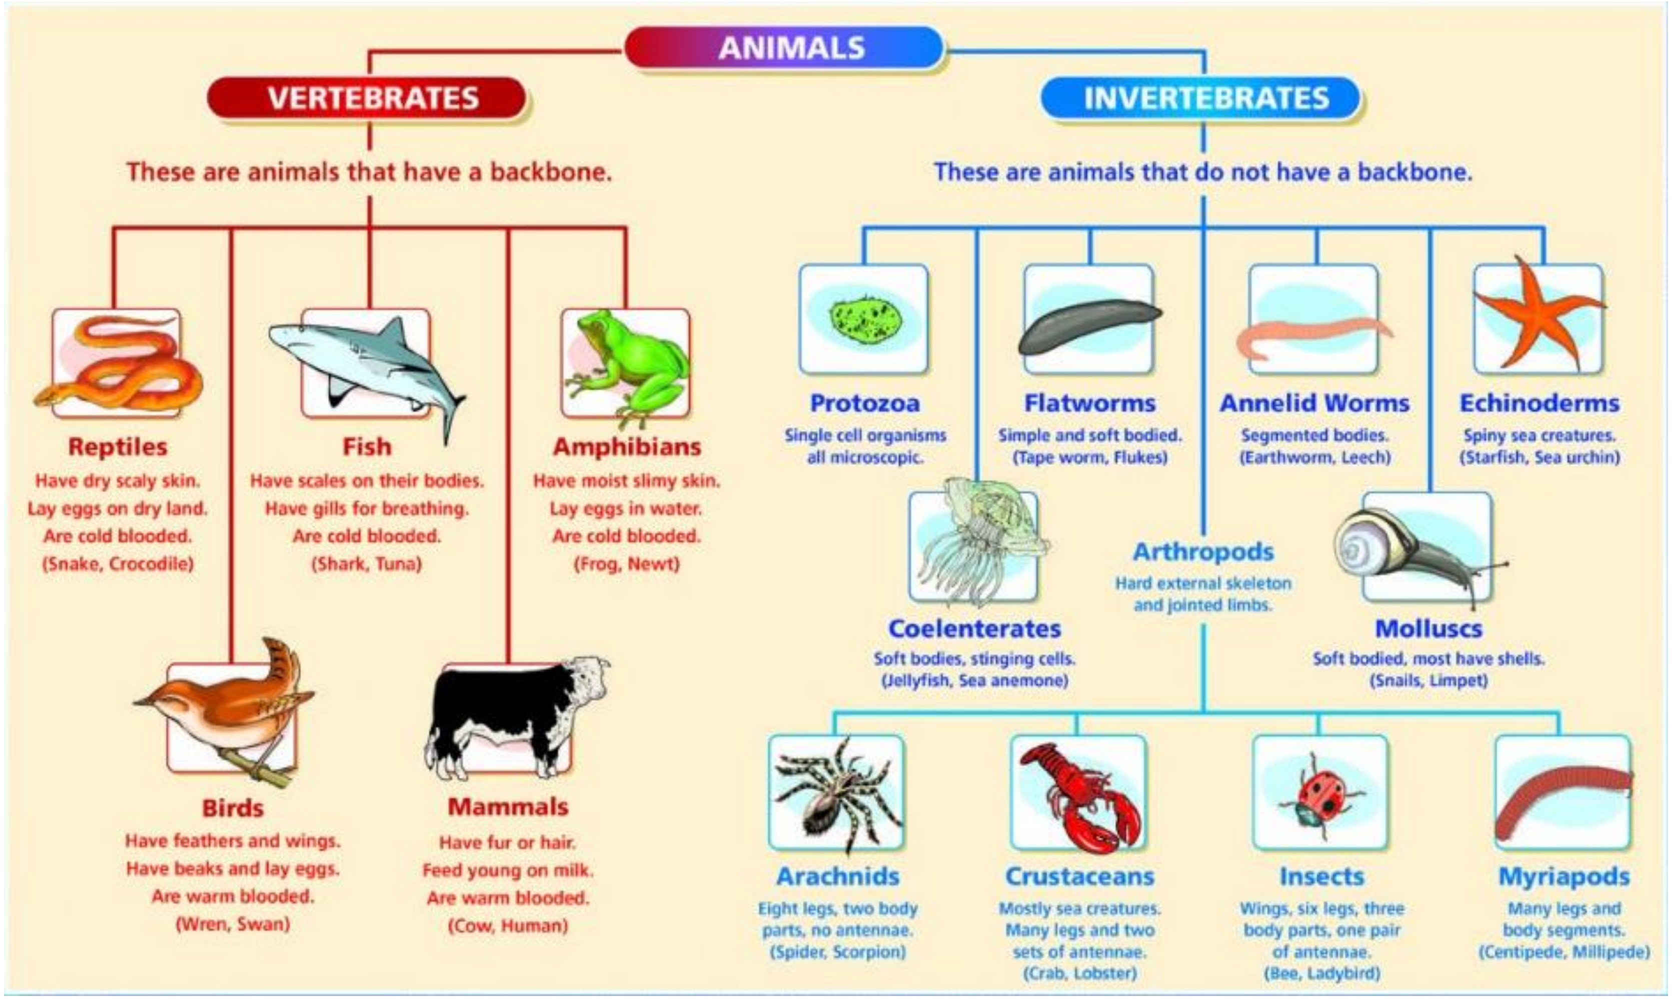

Almost all!

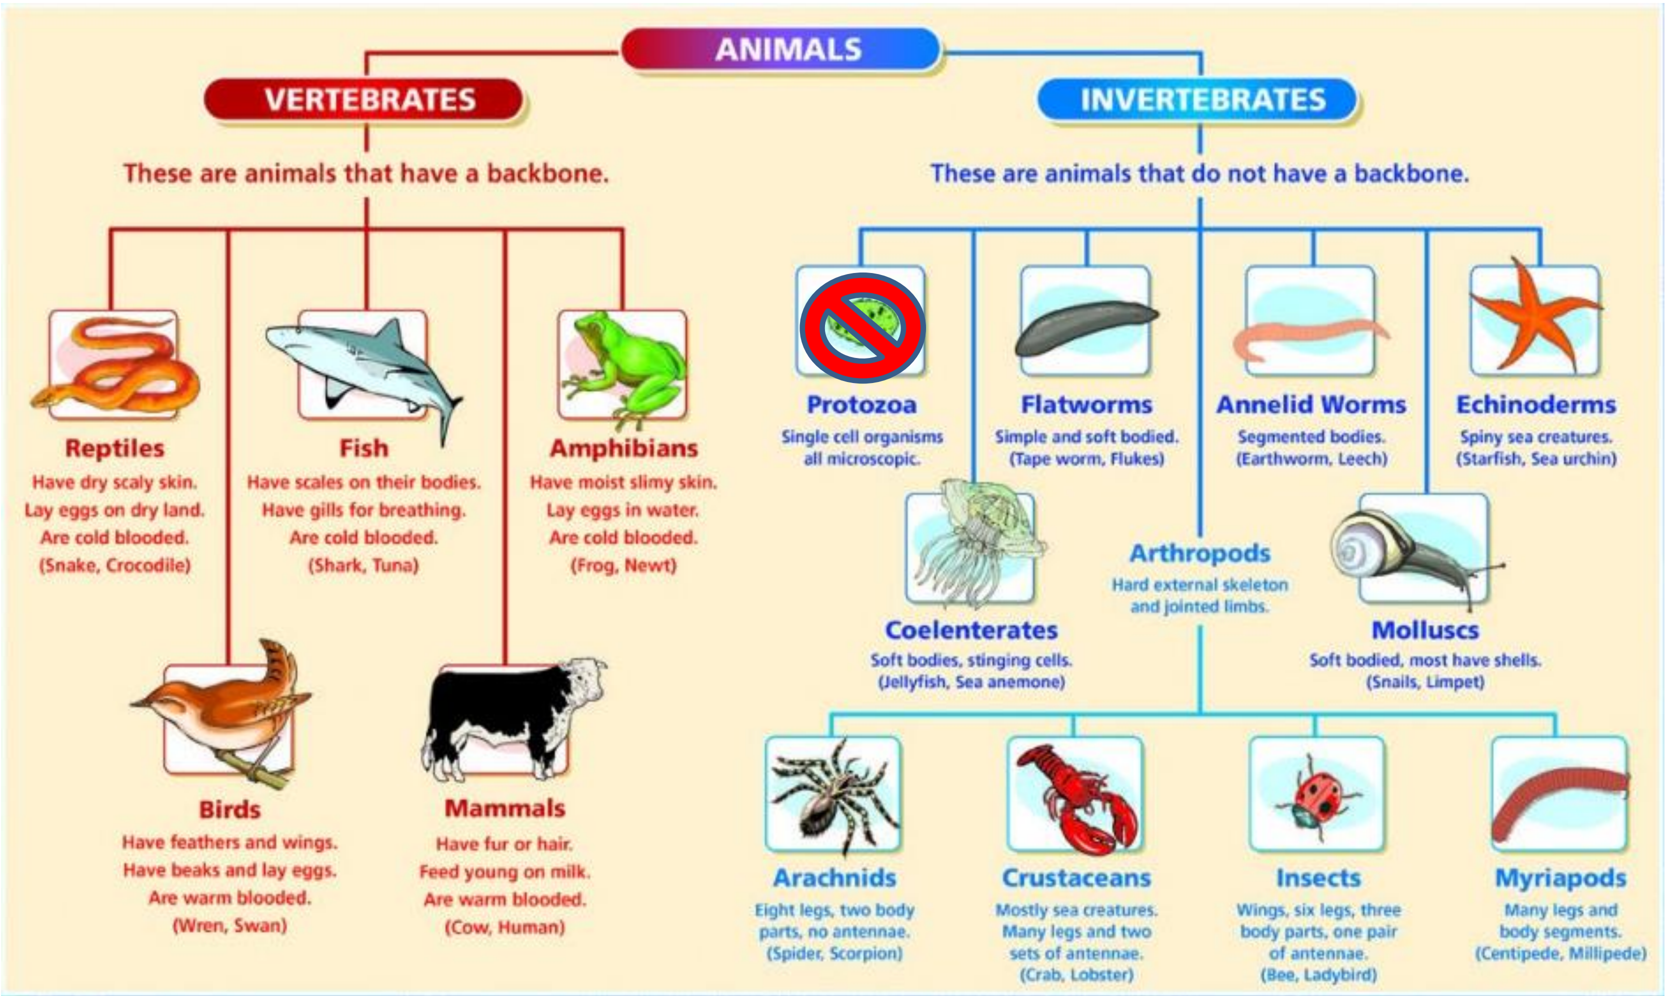
ANIMALS
VERTEBRATES
These are animals that have a backbone.
INVERTEBRATES
These are animals that do not have a backbone.
Reptiles
Have dry scaly skin.
Lay eggs on dry land.
Are cold blooded.
(Snake, Crocodile)
Fish
Have scales on their bodies.
Have gills for breathing.
Are cold blooded.
(Shark, Tuna)
Amphibians
Have moist slimy skin.
Lay eggs in water.
Are cold blooded.
(Frog, Newt)
Birds
Have feathers and wings.
Have beaks and lay eggs.
Are warm blooded.
(Wren, Swan)
Mammals
Have fur or hair.
Feed young on milk.
Are warm blooded.
(Cow, Human)
Protozoa
Single cell organisms
all microscopic.
Flatworms
Simple and soft bodied.
(Tape worm, Flukes)
Annelid Worms
Segmented bodies.
(Earthworm, Leech)
Echinoderms
Spiny sea creatures.
(Starfish, Sea urchin)
Coelenterates
Soft bodies, stinging cells.
(Jellyfish, Sea anemone)
Arthropods
Hard external skeleton
and jointed limbs.
Molluscs
Soft bodied, most have shells.
(Snails, Limpet)
Arachnids
Eight legs, two body
parts, no antennae.
(Spider, Scorpion)
Crustaceans
Mostly sea creatures.
Many legs and two
sets of antennae.
(Crab, Lobster)
Insects
Wings, six legs, three
body parts, one pair
of antennae.
(Bee, Ladybird)
Myriapods
Many legs and
body segments.
(Centipede, Millipede)

This suggests that

- Animals have individual anatomies:
  - needs
  - sensors
  - motors
  - memory

- They also have generic mechanisms for learning and decision-making.

Model the mechanisms, not the biological implementations!

Generic animat idea

- Let us build an artificial animal with a configurable initial "brain" and generic mechanisms for learning and decision-making!

Dynamic graphs

- We use dynamic graphs for modeling "brains"

- A *dynamic graph* consists of
  - A set of nodes that can be of type SENSOR, MOTOR, AND, OR, NOT, SEQ, ACTION, or NEED. Each node also has a unique name.
  - Labeled edges

A dynamic graph

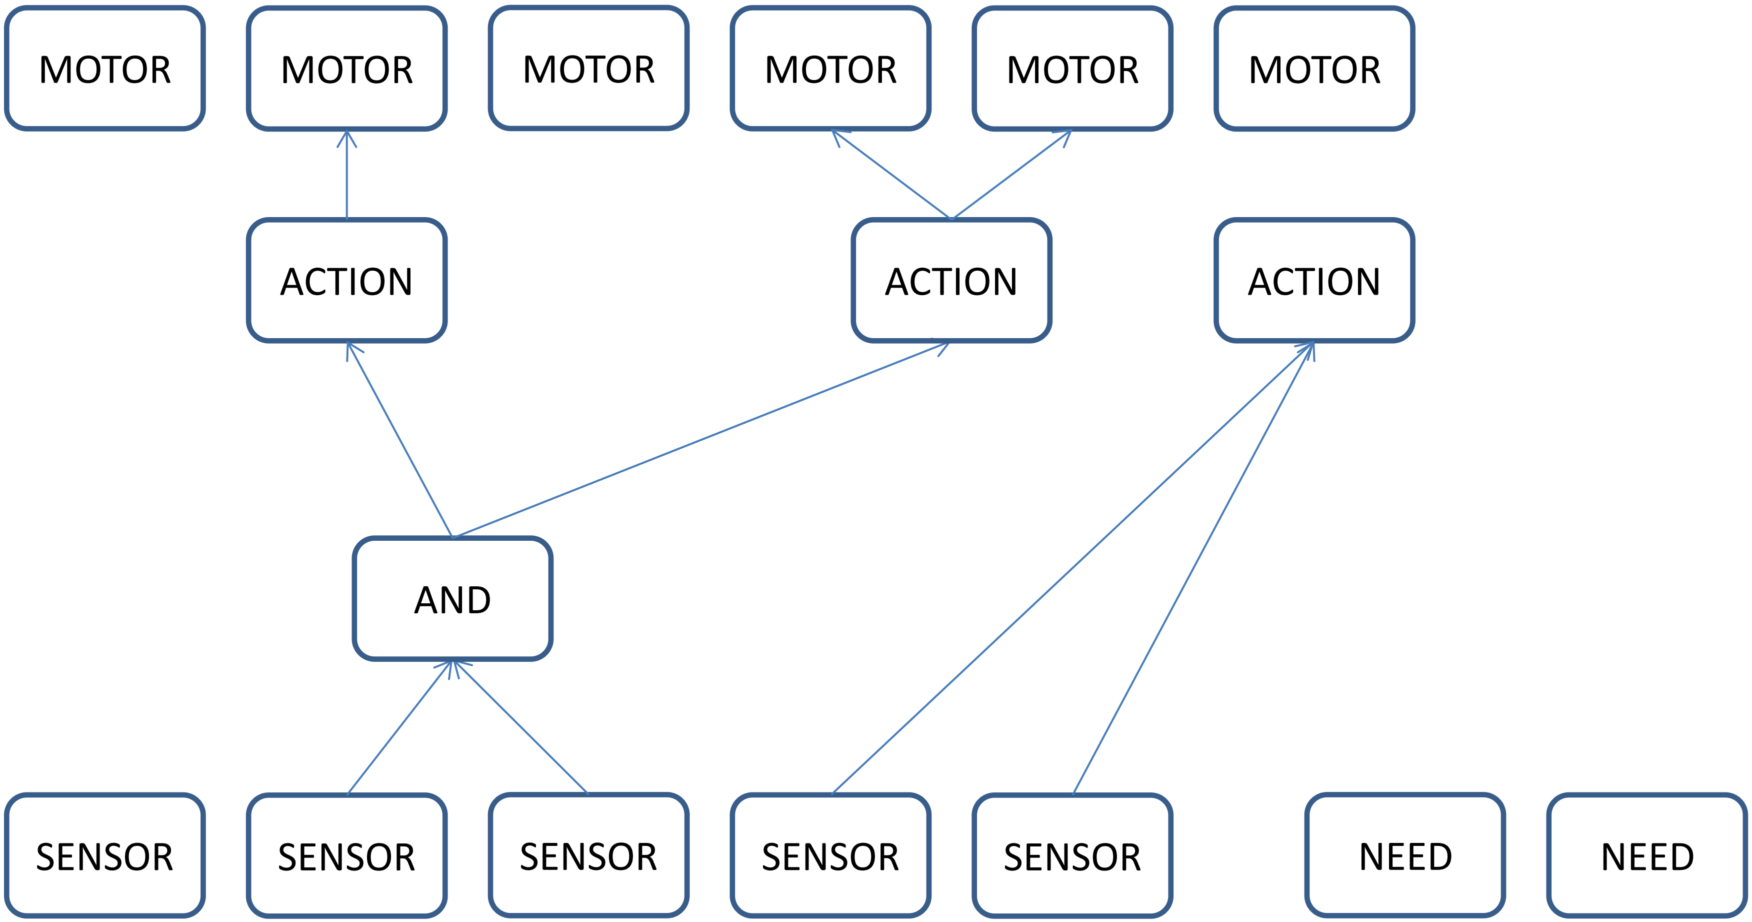

The same graph with names shown

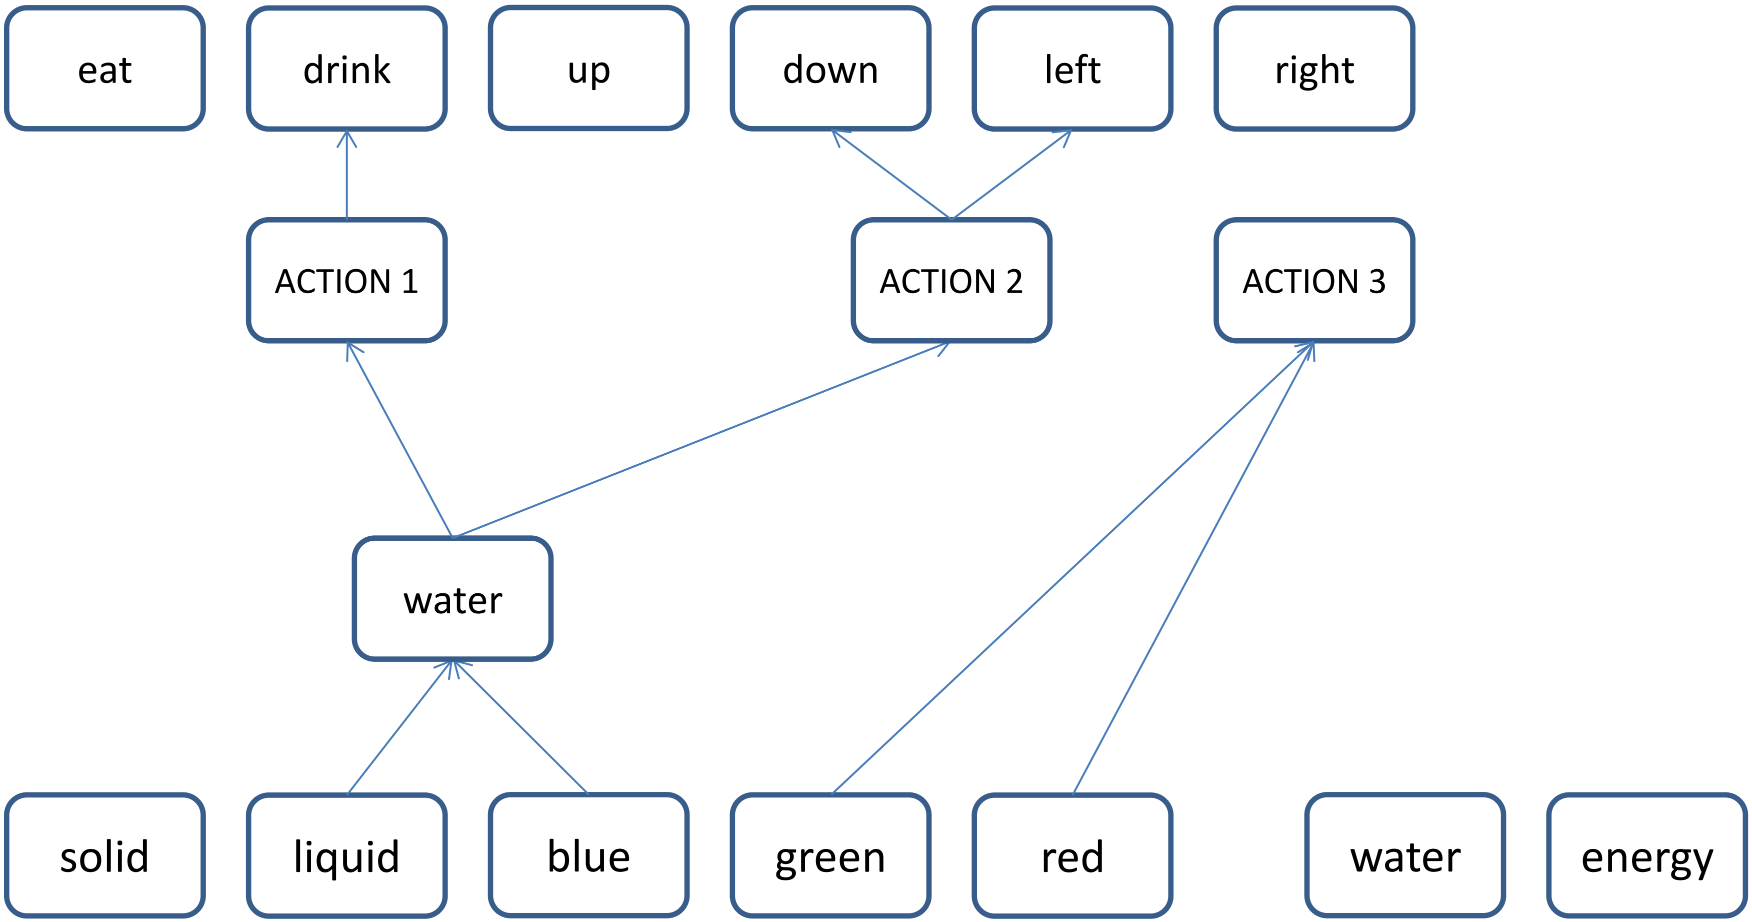

Input

- SENSOR nodes receive inputs from {0,1}
- NEED nodes receive input from [0,1].

Activity propagation

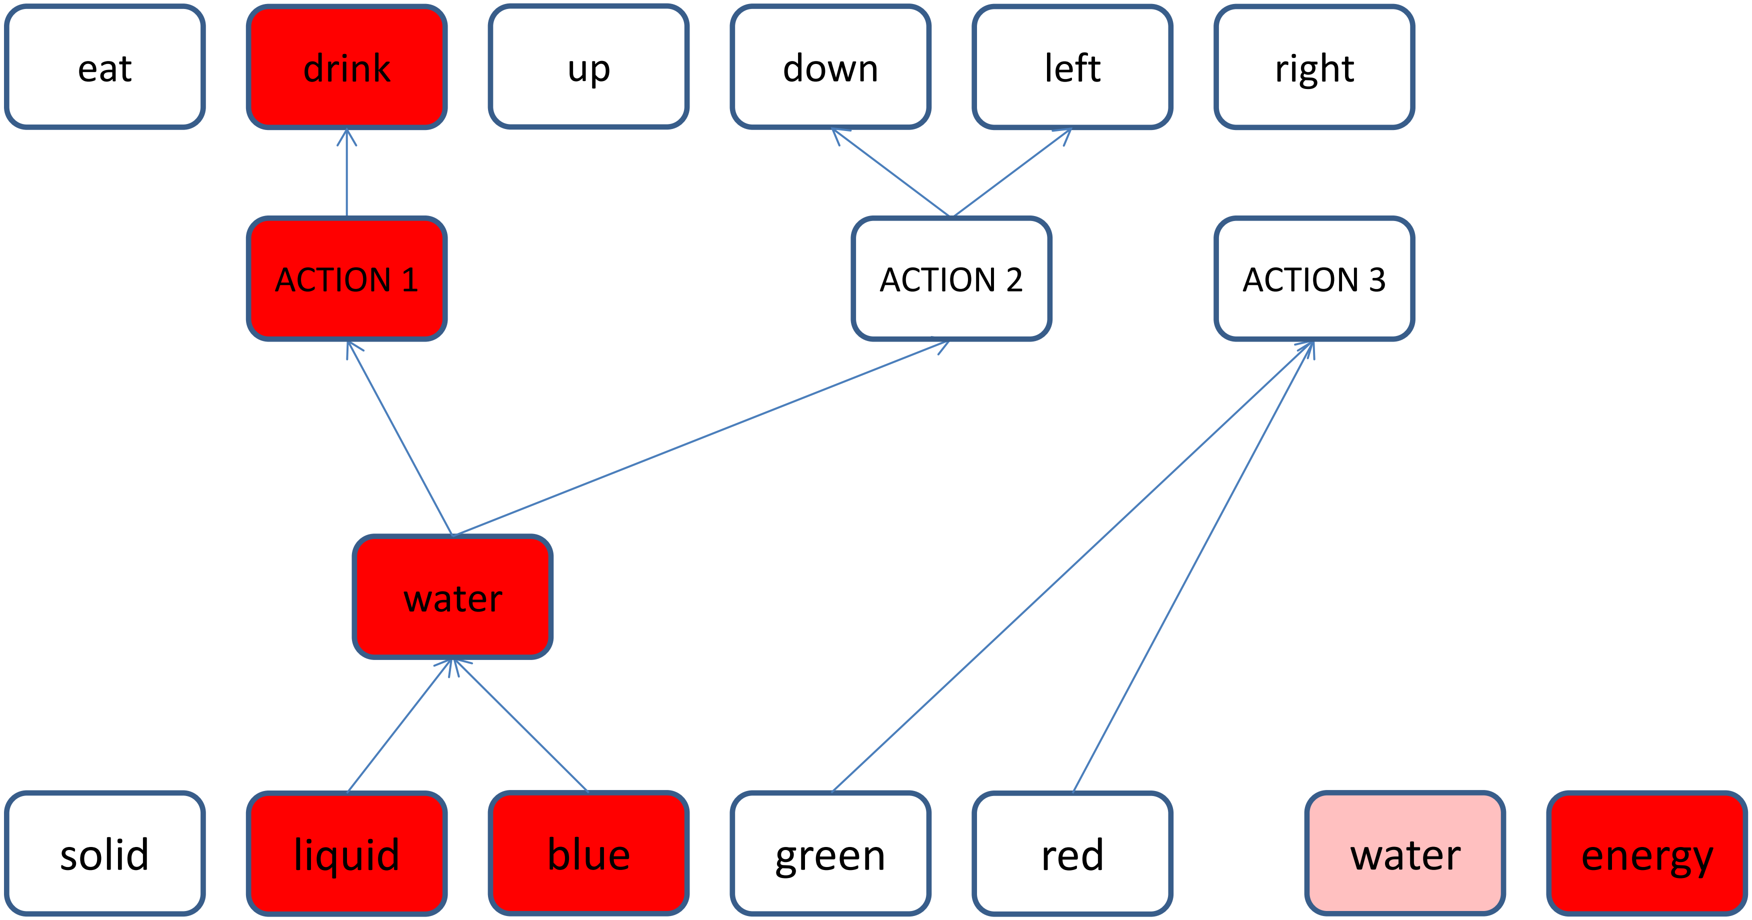

Worlds

- A *block* consists of
  - A block type, e.g. sand, water, or grass
  - A position in space
- A *world* is a set of blocks.

A world

Animats

- An *animat* consists of
  - A dynamic graph G (the animat's "brain")
  - An activity on G
  - A position in a world

Animats in a world

1. Learning what to eat and drink

Sheep at the start

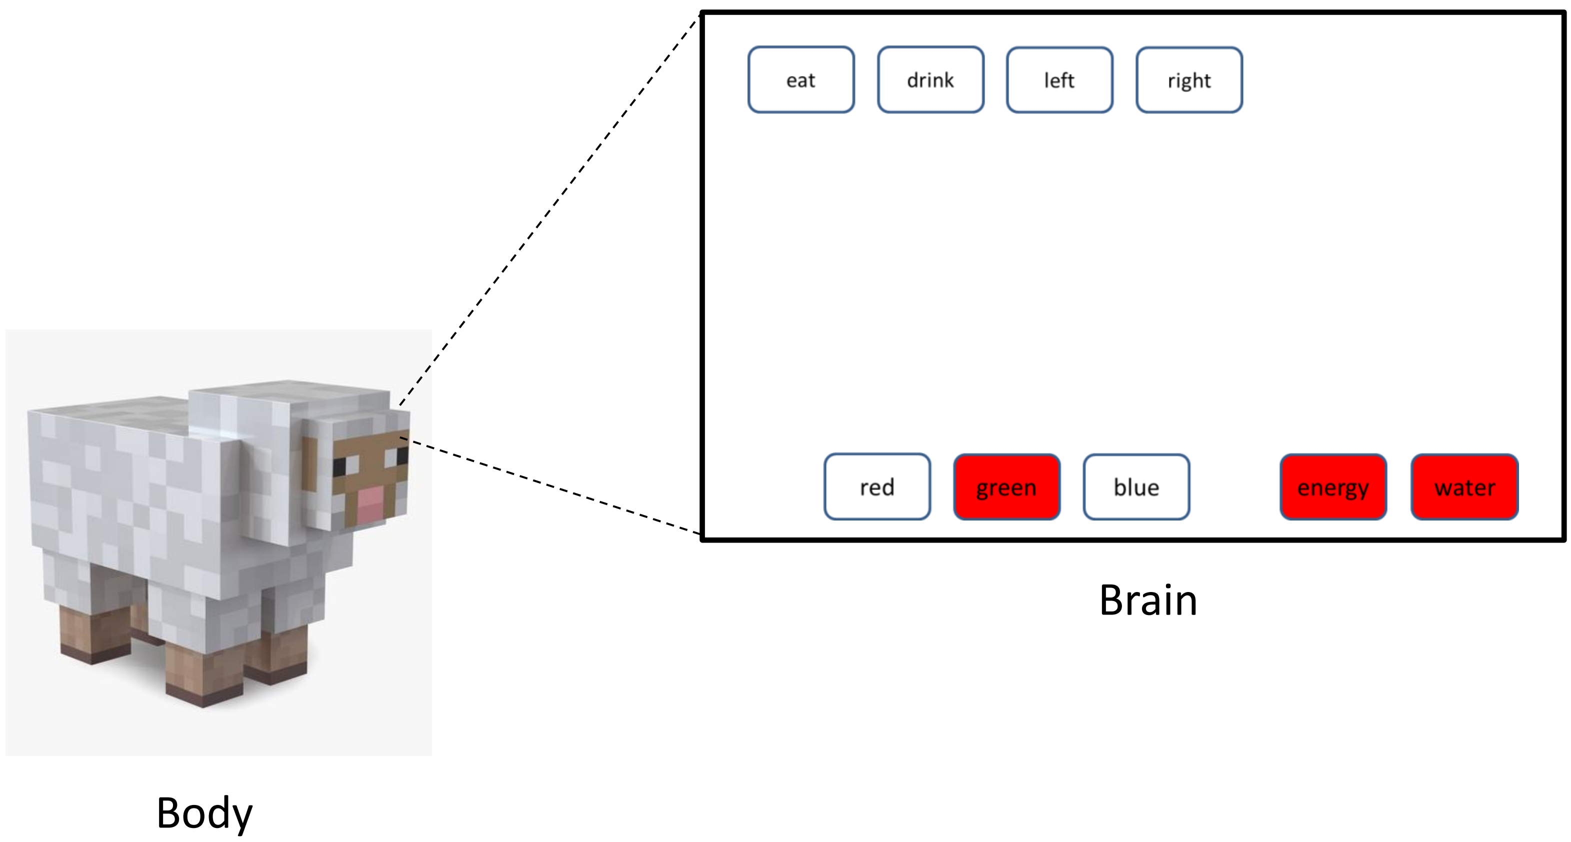

Sheep world

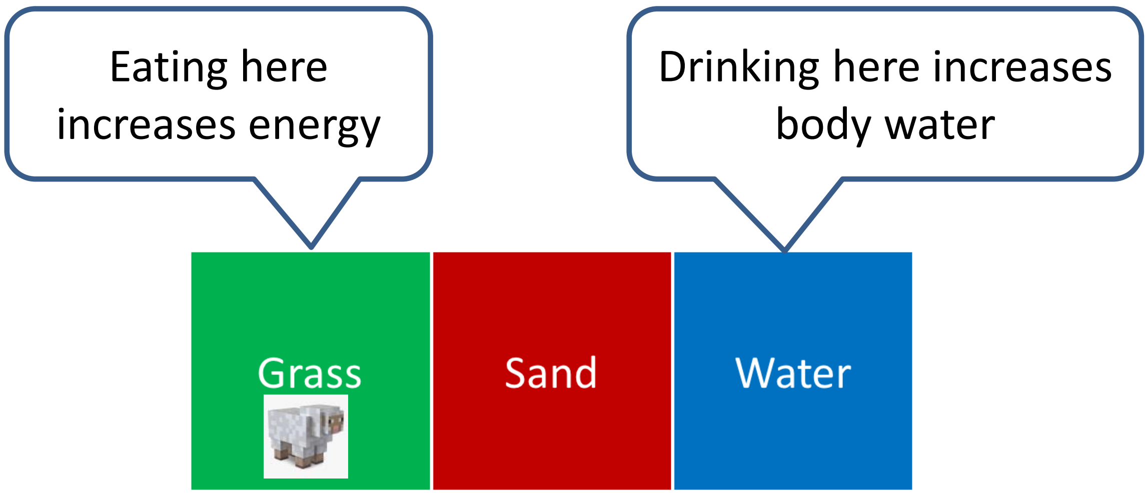

All other actions lead to decreased levels of energy and body water (cost of metabolism).
To keep its needs satisfied it must commute between the Grass and the Water.

Sheep after convergence

Body

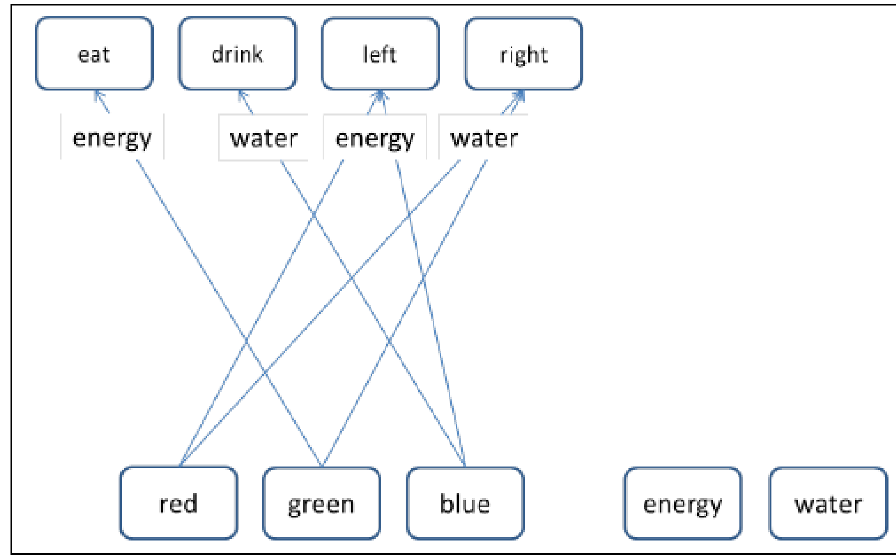

Brain

The labels on the edges indicate the policy depending on the most urgent need.

Vitality

- Vitality
  - $\mathrm{v(t)} = \min$ {water_level(t), energy_level(t)}

- Death
  - If $\mathrm{v(t)} = 0$

Performance of the sheep

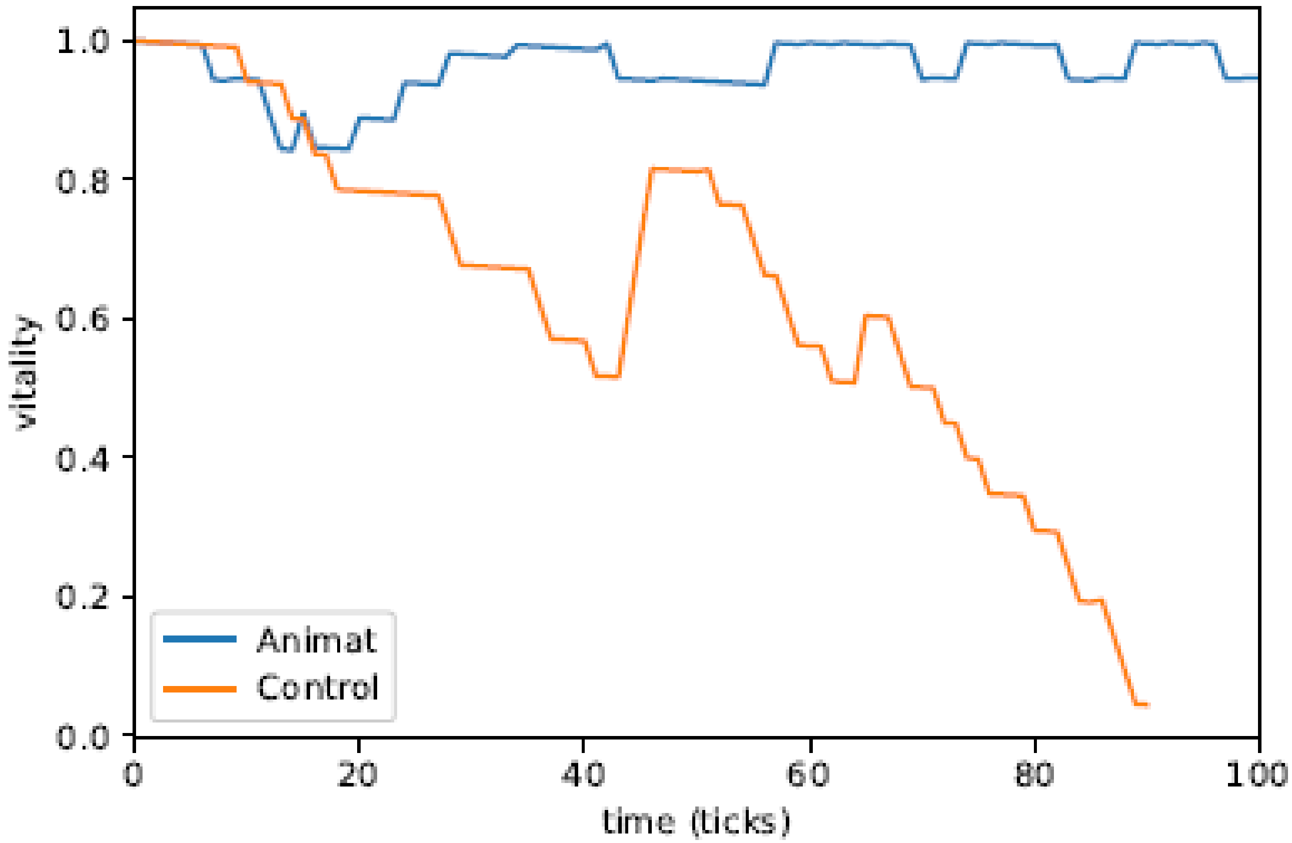

Blue curve: performance of our sheep.
Red curve: performance of a sheep that acts randomly.

2. Learning to move

Toad

To move forward (crawl) the toad must keep moving one hind leg after the other.

Toad at the start

EXTEND LEFT LEG

EXTEND RIGHT LEG

EXTEND BOTH

LEFT LEG RELAXED

RIGHT LEG RELAXED

ENERGY

Toad world

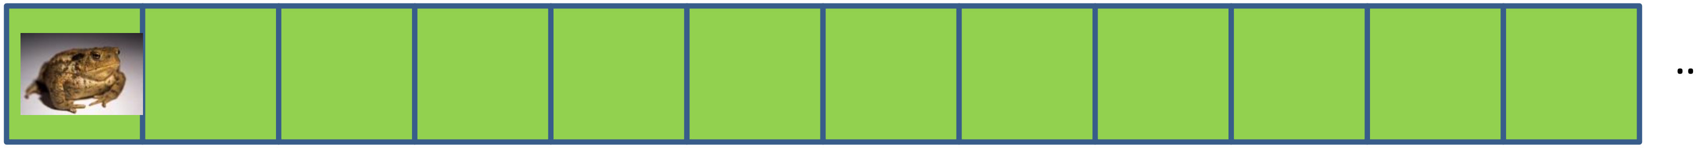

The toad gets a small fly at each block. It needs to keep moving forward to survive.

Toad after convergence

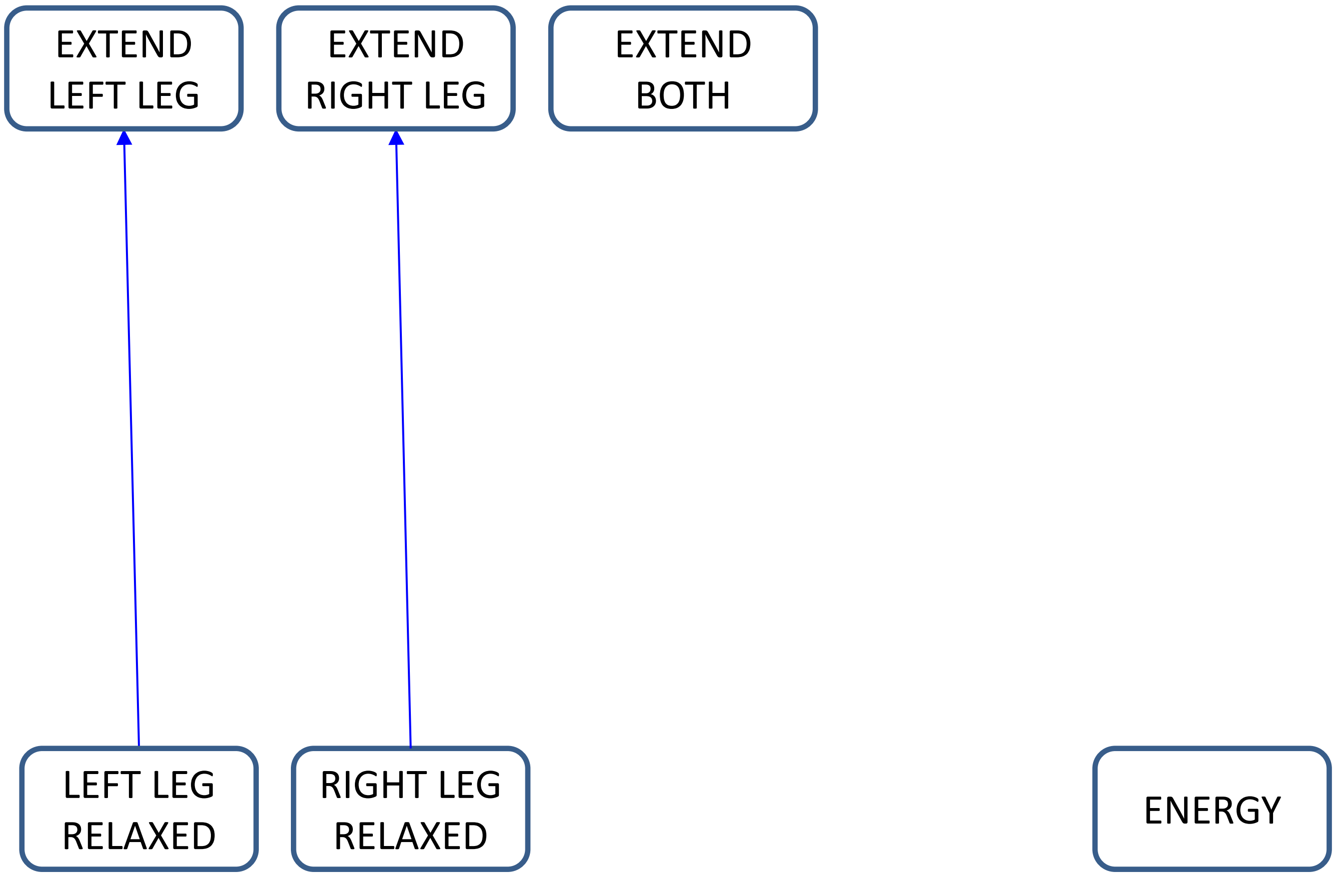

Performance of the toad

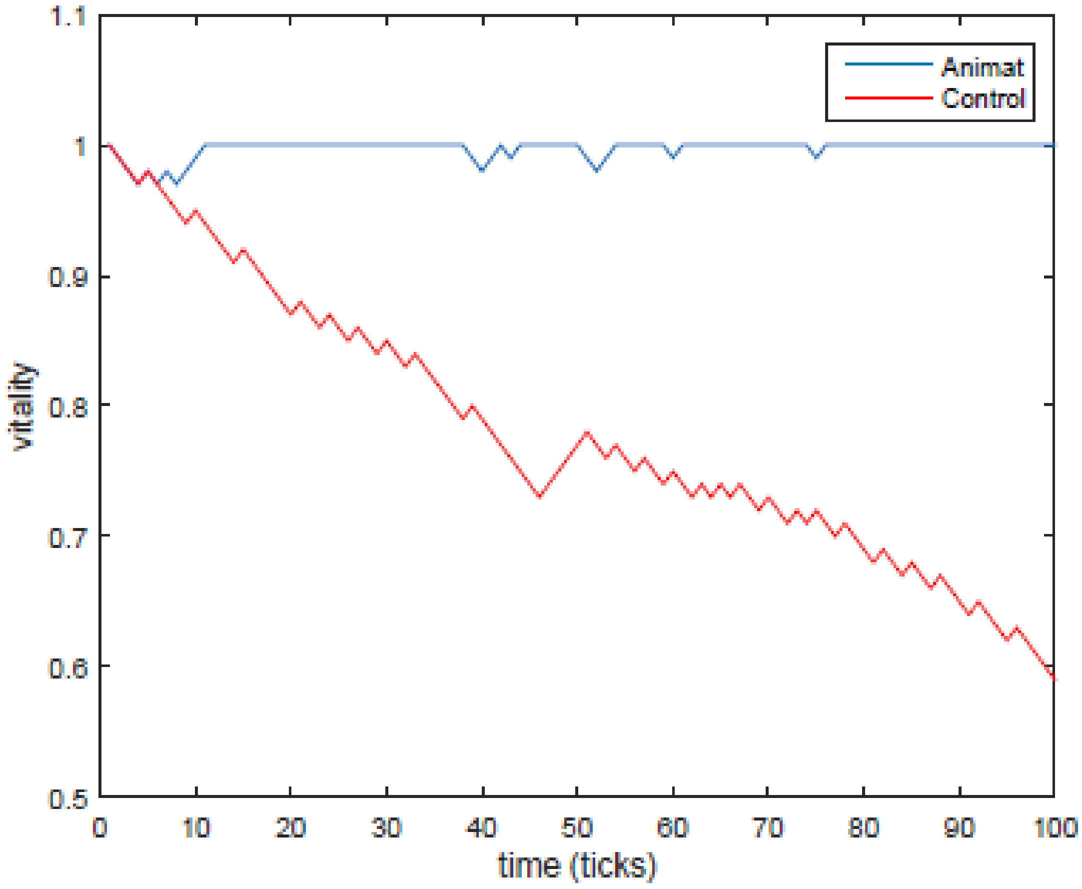

Blue curve: our toad (with proprioception)
Red curve: a similar toad but without proprioception.

Frog

To move forward (jump) it must keep moving both hind leg simultaneously.

Frog after convergence

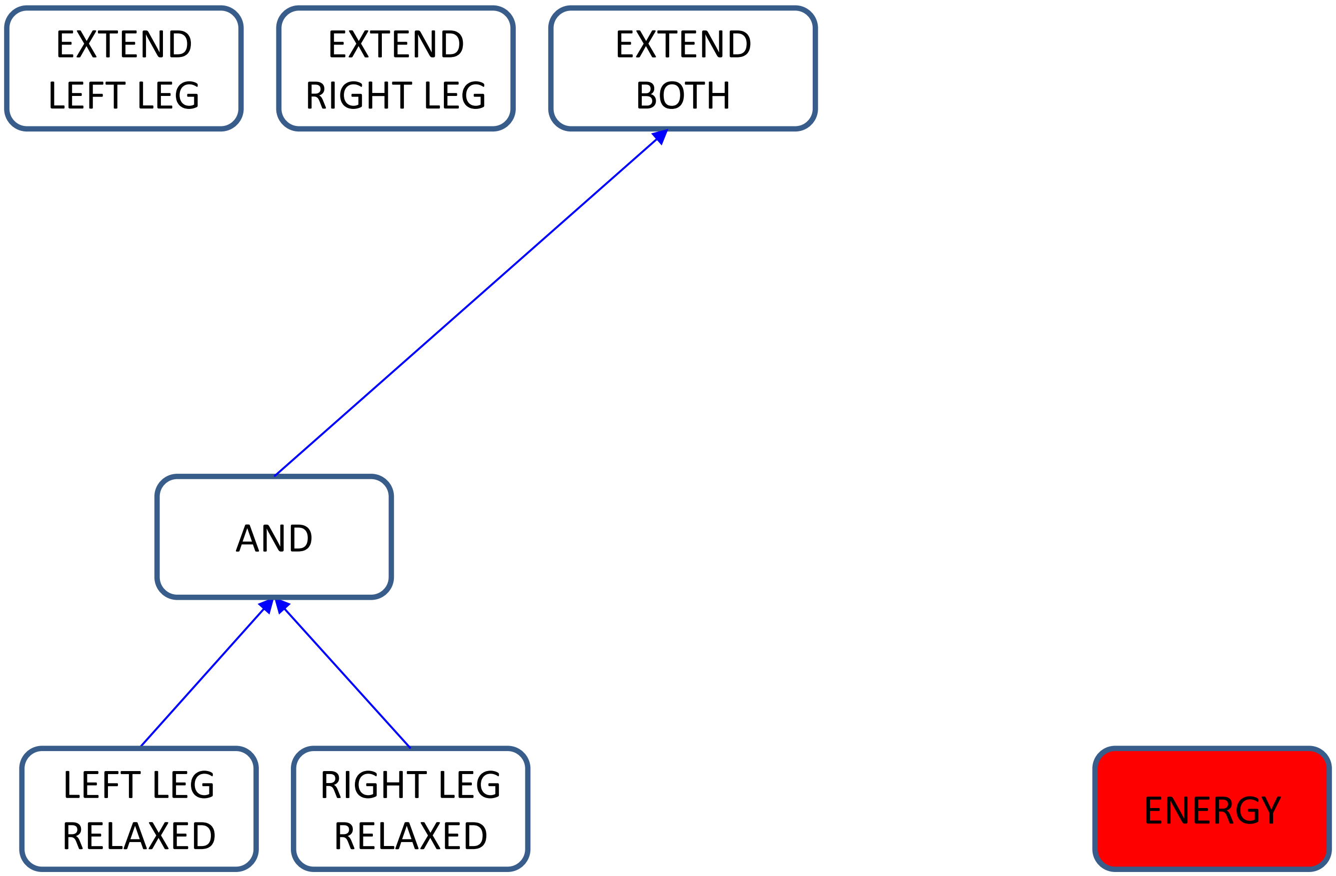

A new AND node is added automatically. One-shot learning triggered by need-related surprise.

3. Learning to navigate

Bee

Bee world at the start

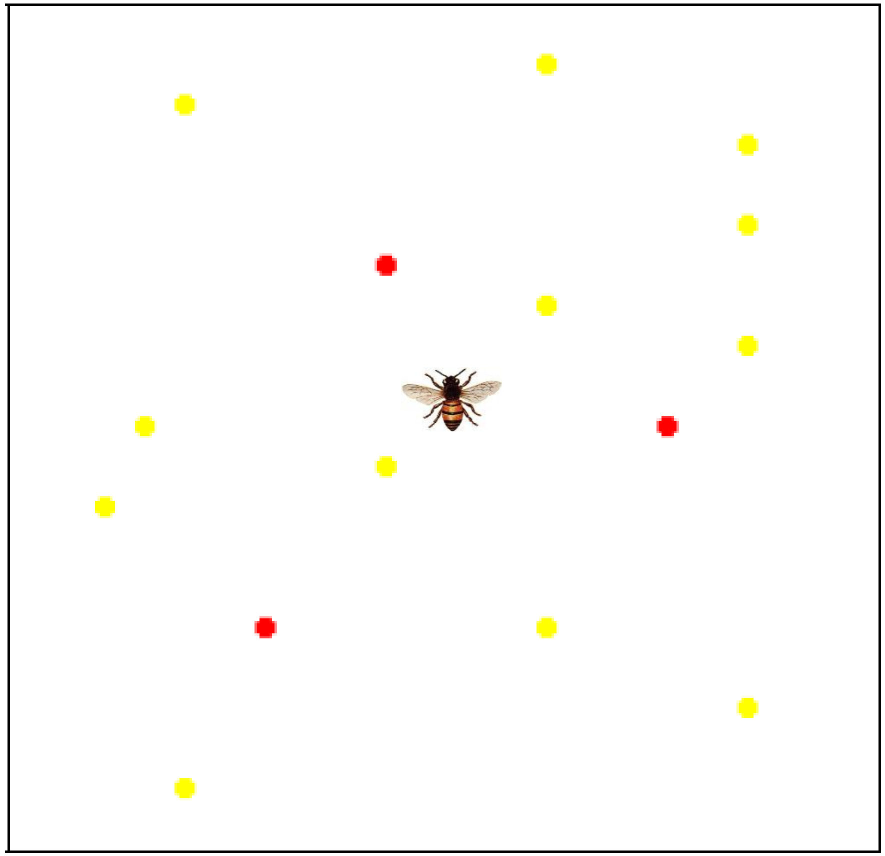

A 20 *x* 20 blocks world with a torus topology. Yellow blocks represent edible flowers.

Bee at the start

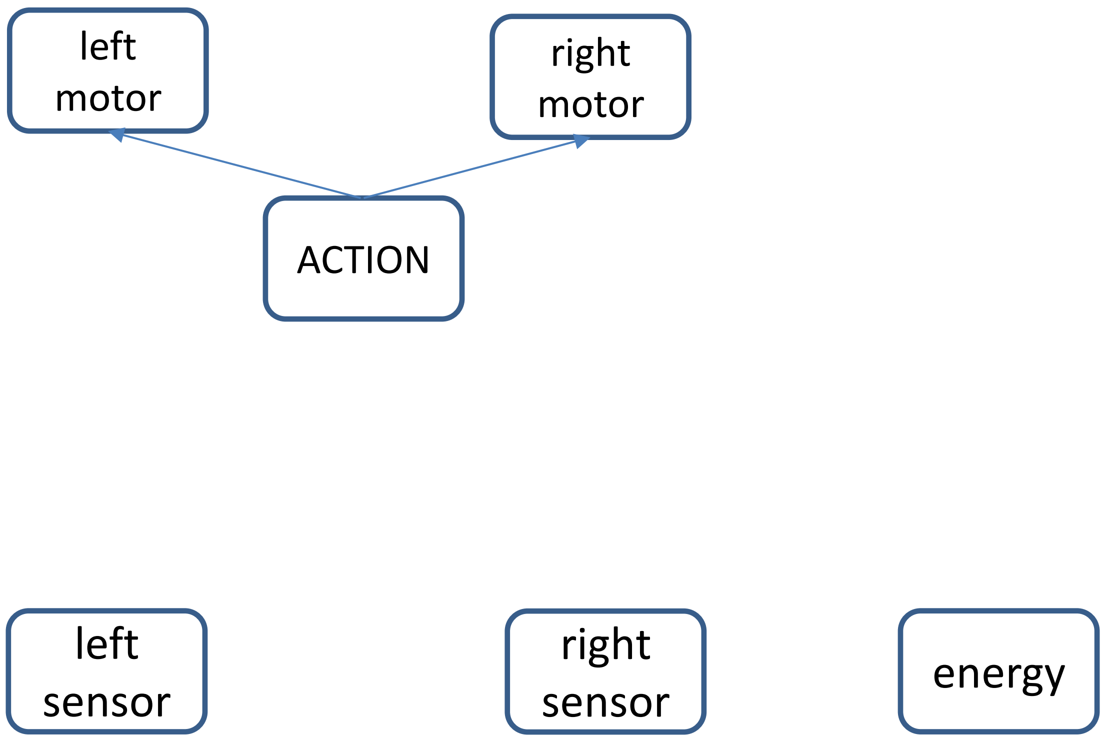

The sensors are sensitive to the gradient of the flower smell.

Bee world during foraging

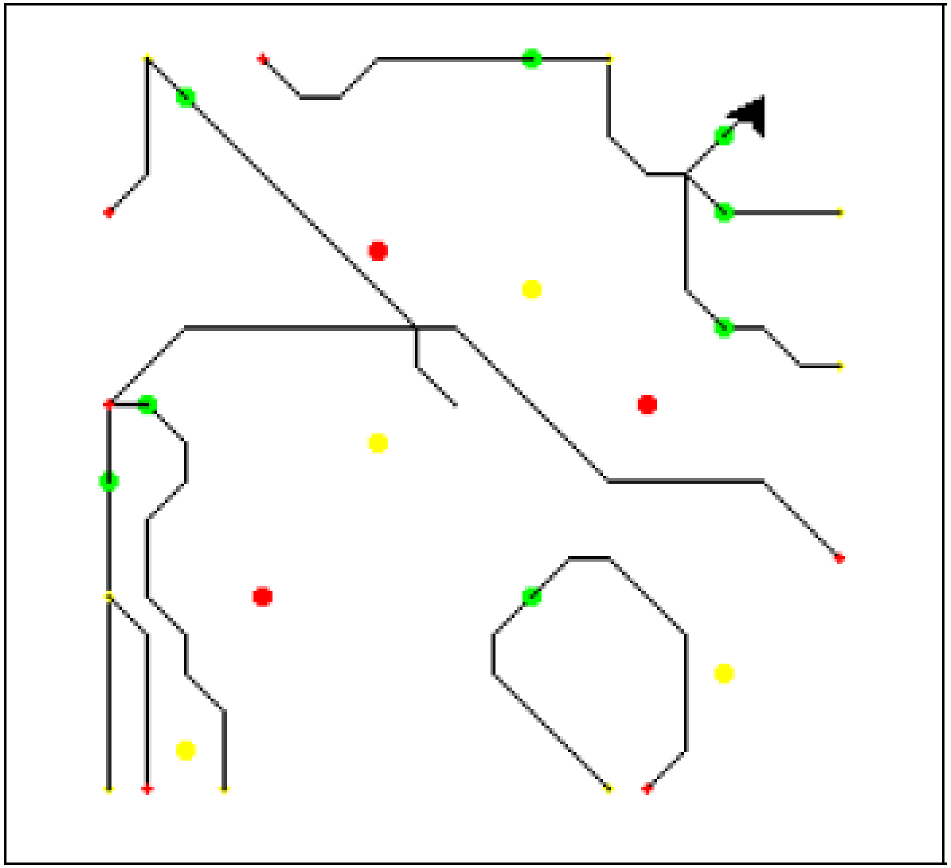

Traces of the bee's movements are shown.

Bee after convergence

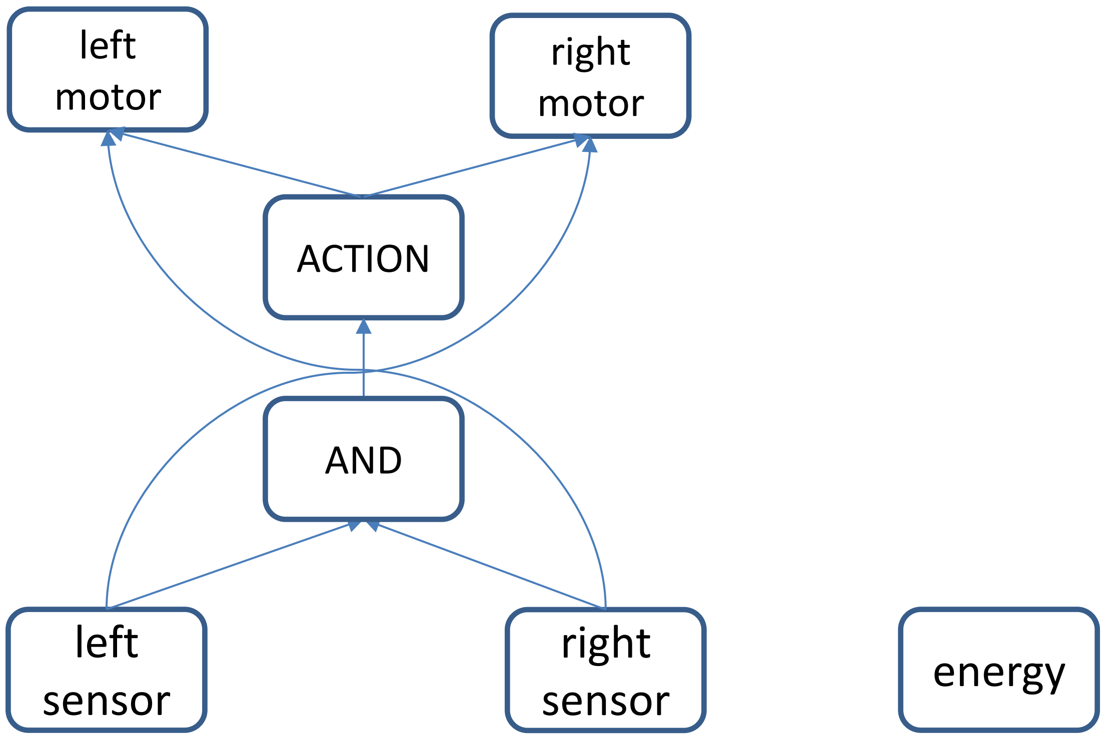

A Braitenberg vehicle of type II (steering towards the gradient) is formed automatically.

Performance of the bee

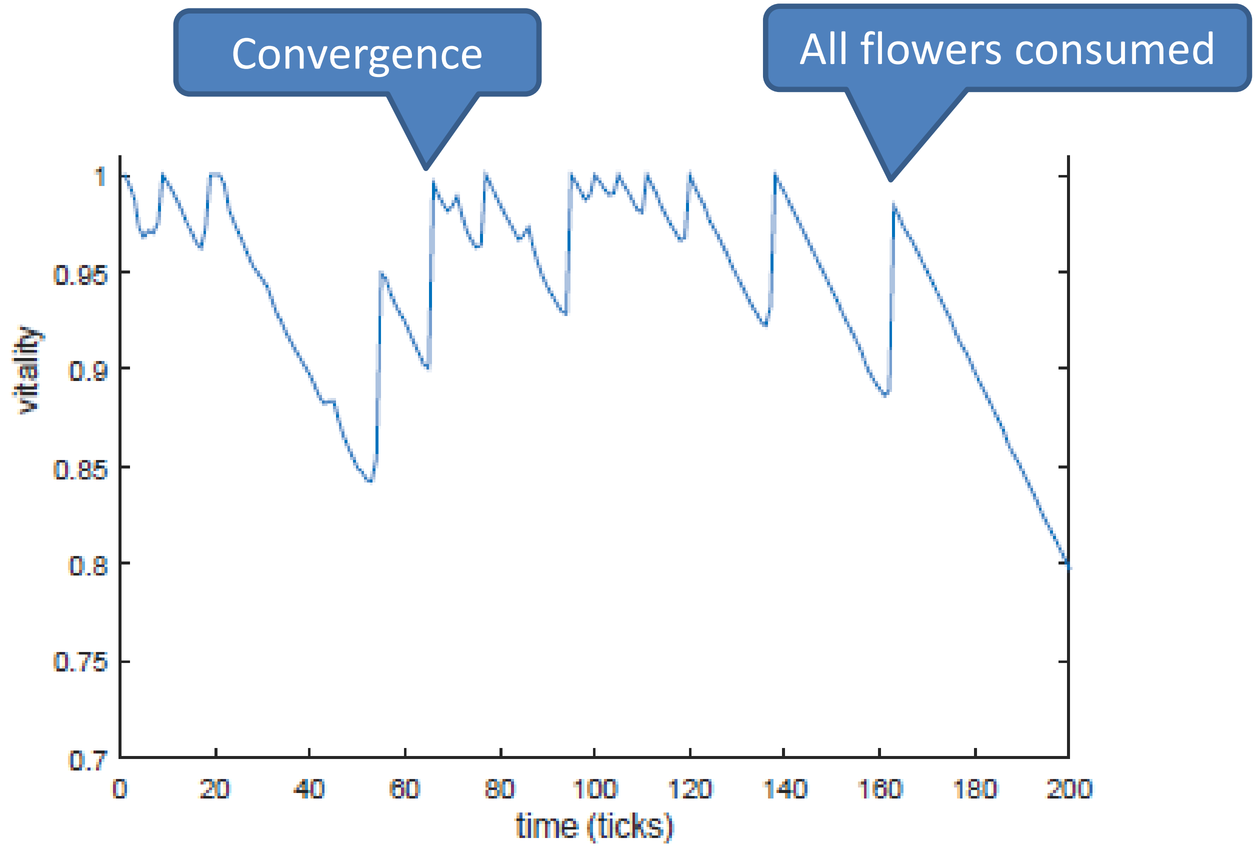

Conclusion

- We presented the Generic Animat: an autonomous system that is able to start with an arbitrary "brain" and adapt to an arbitrary world
- We saw how the Generic Animat can control a sheep, toad, frog, or bee animat and make them adapt to different worlds
- The Generic Animat is intelligent in van Heerden's sense
- The dynamic network topology is key: it seems to outperform all static topologies
- More examples in the AGI-17 proceedings
- Code at github.com/strannegard/ecosystem