

Human-like artificial creatures

3. Norns & neural networks

Cyril Brom

Faculty of Mathematics and Physics
Charles University in Prague
brom@ksvi.mff.cuni.cz

(c) 2/2006

Outline

1. Gentle introduction to Norns
 - the architecture
 - the body
 - the breeding algorithm
2. Gentle introduction to neural networks
3. A network of a norn's brain
4. Possibilities and limitations

⇒ norns

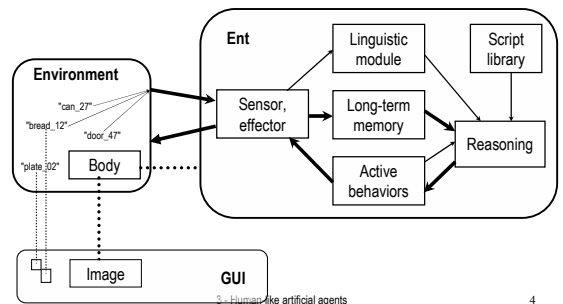
Ad norns...



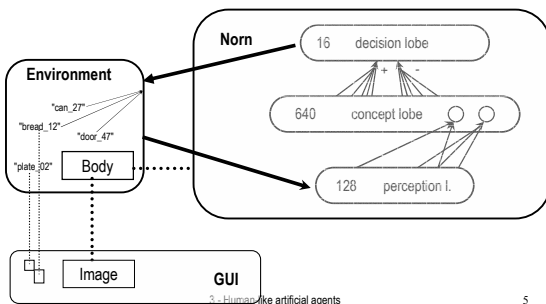
- They live in an artificial environment similar to natural world
- User can observe them and interact with them in a 2,5D GUI
- They are controlled by a neural network comprising about 1000 neurons.
- They speak using "verb-object" sentences
- They are equipped with a body
 - drives: hunger, thirst,...
- They grow and breed through sexual reproduction

1996, Grand, Clif (c)

Overall architecture of a cognitive beast

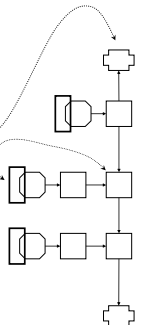


Overall architecture of a norn



The body

- About 20 internal drives
 - fear, boredom, anger, hunger,...
- Biochemistry
 - chemicals (concentrations of H₂O, glucose,...)
 - emitters (emit chemicals)
 - receptors (monitors concentrations)
 - reactions: glucose ⇒ CO₂ + H₂O + energy
- Chemicals are produced by emitters
 - by acting in the world
 - by internal emitters
- Chemicals reduce or increase drives
- Biochemistry is defined genetically in partial



Evolutionary algorithms

- Evolutionary algorithms:
 - an umbrella term used to describe computer-based problem solving systems which use computational models of some of the known mechanisms of evolution
- 1975: Genetic Algorithms:
 - John Holland: "Adaptation in Natural and Artificial Systems"
 - a model of machine learning which derives its behavior from a metaphor of some of the mechanisms of evolution in nature
 - abstraction at the level of individuals
 - a population of individuals represented by chromosomes, in essence a set of character strings... the individuals in the population then go through a process of simulated "evolution"
 - strings = "genome"
 - encoding solutions to the "genomes"
 - fix-length, linear (typically)

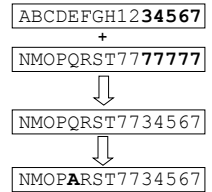
[Hitchhiker's guide]

3 - Human-like artificial agents

7

Genetic algorithms

- A population
- Generations
 - parents vs. offspring
- Operators
 - crossover, mutation, selection
 - fitness
 - "take the genomes encoding solutions with the better fitness"



3 - Human-like artificial agents

8

PSEUDO CODE
Algorithm GA is

GA

```
// start with an initial time
t := 0;

// initialize a usually random population of individuals
initpopulation P (t);

// evaluate fitness of all initial individuals of population
evaluate P (t);

// test for termination criterion (time, fitness, etc.)
while not done do
  // increase the time counter
  t := t + 1;

  // select a sub-population for offspring production
  P' := selectparents P (t);

  // recombine the "genes" of selected parents
  recombine P' (t);

  // perturb the mated population stochastically
  mutate P' (t);

  // evaluate its new fitness
  evaluate P' (t);

  // select the survivors from actual fitness
  P := survive P,P' (t);
od
end GA.
```

⇒ <http://cs.felk.cvut.cz/~xobitko/ga/>

9

Evolutionary programming

- 1966: Evolutionary Programming
 - Fogel, Owens and Walsh: "Artificial Intelligence Through Simulated Evolution"
 - places emphasis on the behavioral linkage between parents and their offspring
 - abstraction of evolution at the level of reproductive populations (i.e. an individual is a species, not an animal)
 - typically does not use any crossover
 - no constraint on representation
 - e.g. a neural network + mutation operates by perturbing a weight vector
- 1992: Genetic programming:
 - John Koza: "Genetic Programming: On the Programming of Computers by Means of Natural Selection"
 - the extension of the genetic model of learning into the space of programs
 - objects that constitute the population are not fixed-length character strings that encode possible solutions to the problem at hand, they are programs that, when executed, "are" the candidate solutions to the problem
 - usually does not use any mutation
 - parse trees, rather than as lines of code (Lisp)
- Immune programming, Evolutionary strategies...

[Hitchhiker's guide]

10

Genetic algorithms in Creatures

- Haploid chromosome
 - each gene carries the genetic instructions for both sexes
 - time-switching
 - outward characteristics + internal structure
 - "red hair", chemo-receptors, weights in the brain,...
- Cross-over + Mutation

3 - Human-like artificial agents

11

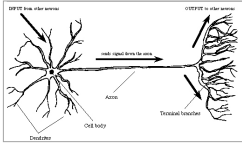
Outline

1. Gentle introduction to norms
2. Neural networks
 - intro
 - a neuron of a norm's brain
 - a topology of the network
 - a perception layer
 - a concept layer
 - dendrits
 - reinforcement learning
3. Possibilities and limitations

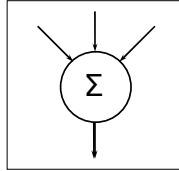
A classical artificial neuron is...

...it is a model of a natural neuron network [McCulloch & Pitts, 1943]

- Oversimplified from the point of view of state-of-art neurobiology
- Every artificial neuron is a function that from an input vector counts an output value
 - weights vs. bias



3 - Human-like artificial agents



13

An artificial neural network is...

...composed of artificial neurons

- Topology:
 - layered, recurrent...
 - input neurons / input layer
 - output neurons / layer
 - a multi-layered perceptron, ADALINE, Hopfield, self-organizing map,...
- Weighted connections



14

Learning in artificial neural networks

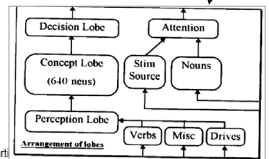
- Supervised
 - with a teacher
 - it gives the correct answer typically
 - e.g. backpropagation
- Reinforcement
 - the environment is the teacher
 - it says whether the action taken has been correct or wrong – reward vs. punishment
 - e.g. q-learning
- Unsupervised
 - without a teacher at all
 - e.g. clustering (based on similarities of input patterns)
- Hebbian principle
 - originally, a hypothesis for how neuronal connections are enforced in mammalian brains [Hebb, 1949]
 - difficult to verify in a mammalian central nervous system



15

A norn's brain

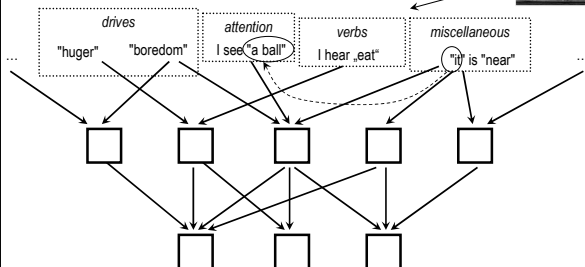
- Composed of several blocks (lobes / layers) of neurons
- Input lobes process internal and external inputs
- Output lobes determine the next action to execute, i.e. "a top-level" task, a script
 - "a water pipe"
- Information spread through the network from input lobes to output lobes



3 - Human-like art

1996, Grand, Clif (c)

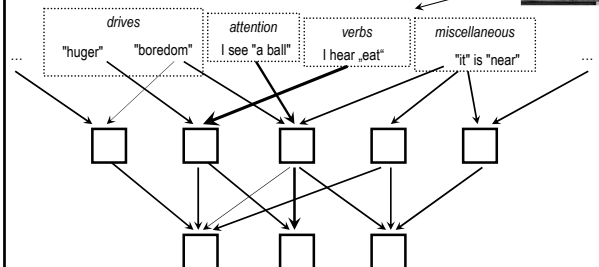
A norn's brain



3 - Human-like artificial agents

17

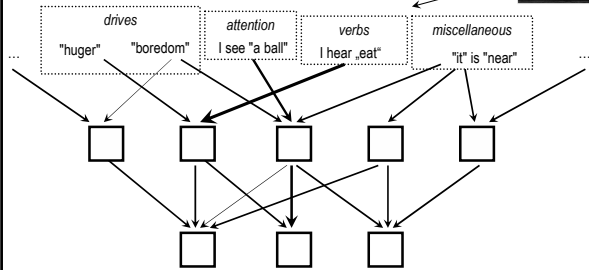
A norn's brain



3 - Human-like artificial agents

18

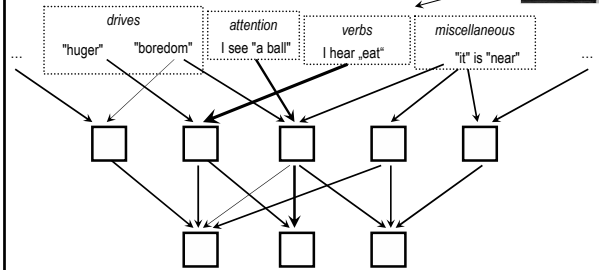
A norn's brain



3 - Human-like artificial agents

19

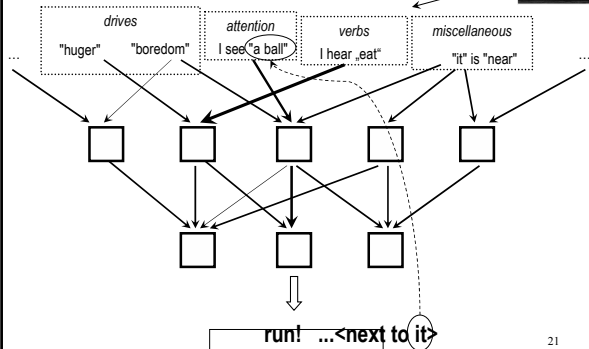
A norn's brain



3 - Human-like artificial agents

20

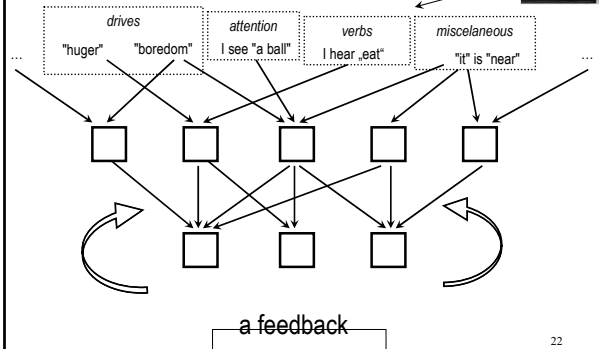
A norn's brain



21

⇒ science kit

A norn's brain

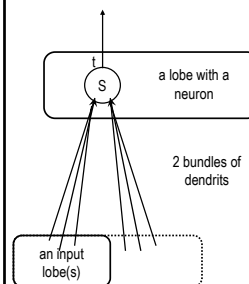


22

Outline

1. Gentle introduction to norns
2. Neural networks
 - intro
 - a neuron of a Norn's brain
 - a topology of the network
 - a perception layer
 - a concept layer
 - dendrits
 - reinforcement learning
3. Possibilities and limitations

A neuron of a norn's brain



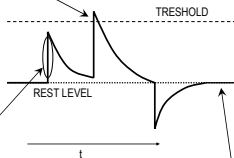
- s – an **internal potential** of the neuron; it goes towards a rest state (a **relaxation rate**)
- s_r – a **rest state**; defined genetically
- t – a **threshold**; an output of the neuron is counted as follows:
 - output = 0 iff $s < t$
 - output = s iff $s \geq t$
- a **neuron SVRule** is a function that computes a value of the potential
 - $s := \text{state PLUS type0}$
 - $s := \text{state PLUS type0 MINUS 1}$
 - $s := \text{anded0}$
- **gain** – a relevance of the neuron

3 - Human-like artificial agents

24

A neuron of the norm's brain

the neuron is excited
(state > threshold)



a stimulus enforces enumeration of the potential (using a neuron SVRule)

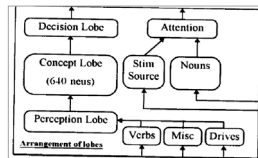
the potential of the neuron returns to its rest level (rest state)

Outline

1. Gentle introduction to norms
2. Neural networks
 - intro
 - a neuron of a Norm's brain
 - a topology of the network
 - a perception layer
 - a concept layer
 - dendrites
 - reinforcement learning
3. Possibilities and limitations

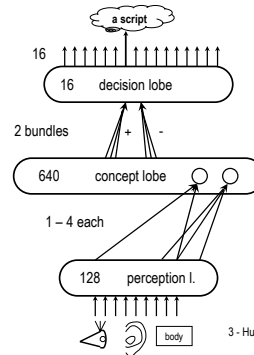
Parts of the brain

- Action selection
- Attention selection



Organization of the brain

Action selection

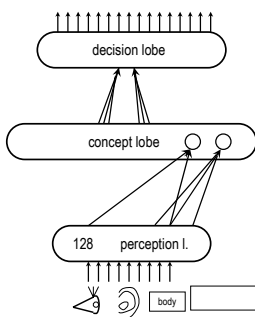


The purpose: to select an action

- **decision lobe** – chooses the best action according to the current situation
- **concept lobe** – remembers interesting/frequent situations (i.e. combinations of perceived symbols)
- **perception lobe** – each neuron represents the presence or absence of a "symbol" in the perceptual system

Organization of the brain

Action selection – an example

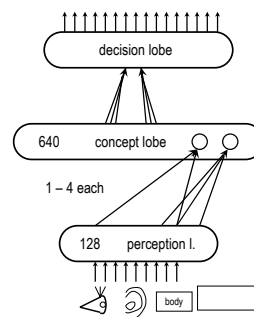


it is perceived –

- sleepiness, tiredness, boredom: high
- seen: a toy + it is "near"
- heard: "play" (verb) + "toy" (object) 29

Organization of the brain

Action selection – an example



useful concepts – these neurons match:

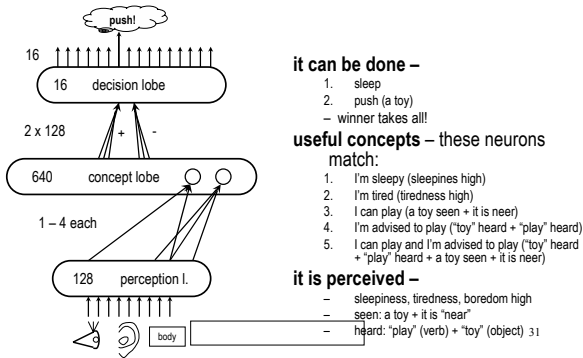
1. I'm sleepy (sleepiness is high)
2. I'm tired (tiredness is high)
3. I can play (a toy is seen + it is near)
4. I'm advised to play ("toy" heard + "play" heard)
5. I can play and I'm advised to play ("toy" heard + "play" heard + a toy seen + it is near)

it is perceived –

- sleepiness, tiredness, boredom high
- seen: a toy + it is "near"
- heard: "play" (verb) + "toy" (object) 30

Organization of the brain

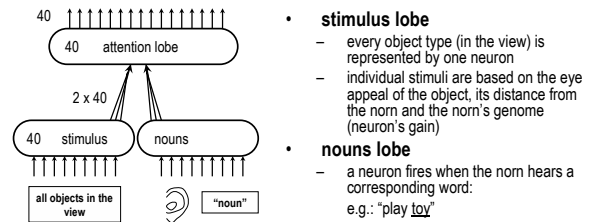
Action selection – an example



Organization of the brain

Attention selection

The purpose: to focus on exactly one object

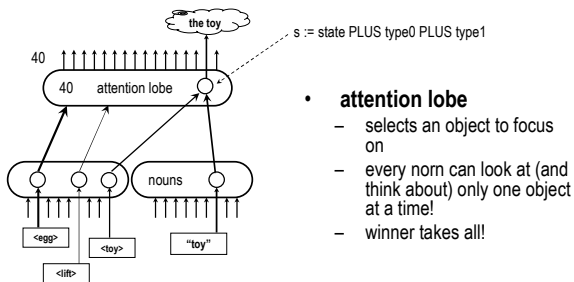


3 - Human-like artificial agents

32

Organization of the brain

Attention selection



3 - Human-like artificial agents

33

Outline

1. Gentle introduction to norms
2. Neural networks
 - intro
 - a neuron of a Norm's brain
 - a topology of the network
 - a perception layer
 - a concept layer
 - dendrits
 - reinforcement learning
3. Possibilities and limitations

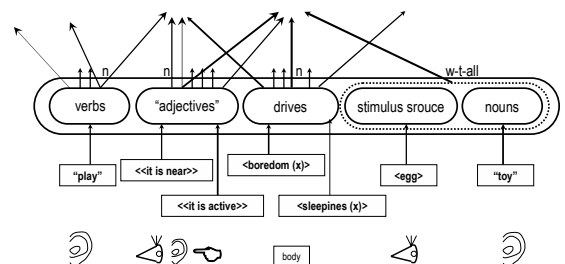
Perception lobe

- It processes visible, audible and internal events
- It contains 4 sub-lobes (technically a copy of their outputs)
- Totally, it contains 128 neurons
 - it means every norm recognizes 128 symbols
- 0-n neurons from each sub-lobe can fire (with one exception)
- All outputs are connected to the concept lobe
 - each neuron in the concept lobe can have 1-4 input dendrits
 - each neuron in the concept lobe can have only one input dendrit from each sub-lobe (with one exception)

3 - Human-like artificial agents

35

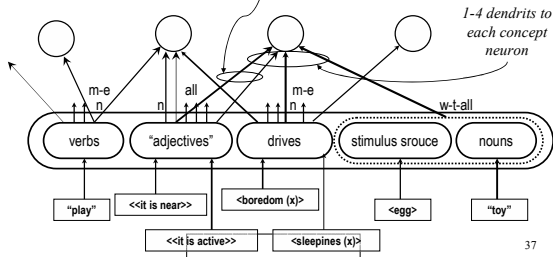
Perception lobe



36

Perception lobe

- Verbs and drives are "mutual-exclusive" lobes, that means at most one dendrit from the verbs/adjectives lobe may feed an individual concept neuron
- The adjective lobe is flagged as "all", that means more neurons from the lobe can participate in one concept



37

Perception lobe

verbs – actions	drives	objects	other ("adjectives")
<ul style="list-style-type: none"> nothing push pull stop come run get drop think sleep walk left walk right eat hit 	<ul style="list-style-type: none"> pain need for pleasure hunger coldness hotness tiredness sleepiness loneliness overcrowdedness fear boredom anger sexdrive injury suffocation thirst stress 	<ul style="list-style-type: none"> current nom user call button water plant egg food drink vendor music animal toy shower weed mover lift nom ... 	<ul style="list-style-type: none"> I've been patted I've been slapped I've bumped into a wall I'm near a wall I'm in the vehicle User has spoken Creatures has spoken Audible event Visible event It is approaching It is retreating It is near me It is an object It is my child It is opposite sex It has hit me ...

38

Outline

- Gentle introduction to norms
- Neural networks
 - intro
 - a neuron of a Norm's brain
 - a topology of the network
 - a perception layer
 - a concept layer
 - dendrits
 - reinforcement learning
- Possibilities and limitations

Concept neuron an example

- A concept could be:
 - "play" + "toy" + "toy" + "boredom (high)"
 - <toy> + <it-is-near-me> + <it-is-active>
 - <nom> + <sex-drive (high)>
 - <nom> + <it-is-opposite-sex> + <sex-drive (high)>
- A concept could be (but is not useful):
 - <hotness (high)> + <toy> + "look" + "egg"
 - this concept will be forgotten
- A concept can not be:
 - <toy> + <egg>
 - "play" + "toy" + "push"

Can both neurons "play" and "push" fire simultaneously?
Why?

3 - Human-like artificial agents

40

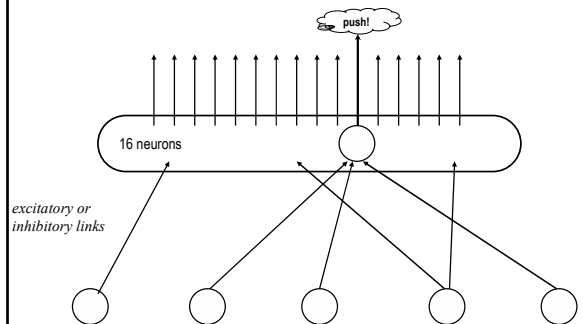
Concept lobe

- Contains 640 concept neurons
- Every concept neuron influences the action selection in a positive or a negative manner
- The initial wiring between the perception lobe and the concept lobe is set genetically with some random noise

3 - Human-like artificial agents

41

Decision lobe



3 - Human-like artificial agents

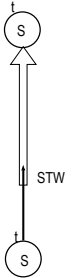
42

Outline

1. Gentle introduction to norms
2. Neural networks
 - intro
 - a neuron of a Norm's brain
 - a topology of the network
 - a perception layer
 - a concept layer
 - **dendrits**
 - **reinforcement learning**
3. Possibilities and limitations

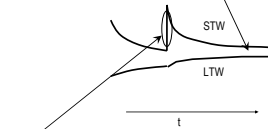
Dendrits

- Everything important happens between neurons
- Main parameters of a dendrit are:
 - **STW** (short-term-weight) – it is not static!
 - **LTW** (long-term-weight)
 - STW/LTW **relaxation rate**
- The output of a neuron is modulated by STW
 - dendrite value = output-of-source-cell * (STW / 255)



Dendrits – STW / LTW

STW goes towards LTW exponentially (relaxation rate) with the LTW being slower



STW is modulated by reinforcement SVRule during a reinforcement episode

LTW computes moving average of many STW disturbances

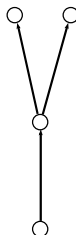
Reinforcement learning

- Hebbian learning – "it could be an interesting concept when both neurons fire – strengthen the weight"
 - a **susceptibility** to reinforcement is increased in such situation (by a **susceptibility SVRule**)
 - the susceptibility goes exponentially to 0 (half-time is defined by a susceptibility SVRule)
- Later, **reward or punishment** can be generated and STW is changed according to it:
 - $STW = STW + (susceptibility / 255) * reinforcement-SVRule$
 - reward and punishment are chemicals
 - "hidden" in the reinforcement SVRule



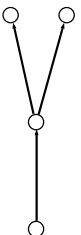
Reinforcement learning

1. Neurons fire, a norm chooses an action



Reinforcement learning

1. Neurons fire, a norm chooses an action
2. All dendrits between all firing p-neurons and firing c-neurons and all firing c-neurons and firing d-neurons are made sensitive



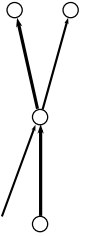
Reinforcement learning

1. Neurons fire, a norm chooses an action
2. All dendrits between all firing p-neurons and firing c-neurons and all firing c-neurons and firing d-neurons are made sensitive
3. Reward or punishment is generated
4. When it "arrives" at a sensitive dendrit, the STW of the dendrit is changed



Reinforcement learning

1. Neurons fire, a norm chooses an action
2. All dendrits between all firing p-neurons and firing c-neurons and all firing c-neurons and firing d-neurons are made sensitive
3. Reward or punishment is generated
4. When it "arrives" at a sensitive dendrit, the STW of the dendrit is changed
5. Susceptibility is being decreased continually, while other neurons are firing



Reward and punishment

How are they generated?

- A set of chemicals is generated inside the body-system when the norm has performed an action (ate food, pushed a toy,...).
- These chemicals increase or decrease a drive and produce the reward chemical or the punishment chemical in the following way:
 - driveraiser => drive + punishment
 - drivereducer + drive => reward
- The meaning of these transformations is:
 - to increase a drive is always bad (e.g. to increase hunger)
 - to decrease a drive is good *iff* the drive is present (e.g. a decrease of "hunger" is good if the norm is hungry)
- Excitatory dendrits are reinforced by reward while inhibitory dendrits are reinforced by punishment

Reinforcement

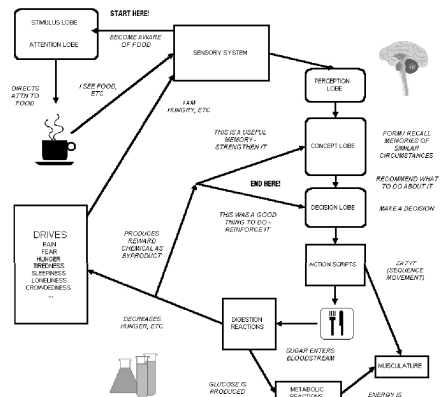
The meaning

- The meaning of an increase (a decrease) of STW of a dendrit between a c-neuron and a d-neuron is:
 - the action you have performed is appropriate (inappropriate) in the given situation
 - it has been appropriate (inappropriate) that you have NOT performed the action in the given situation
- The meaning of an increase of susceptibility of all incoming dendrits of a c-neuron is:
 - this is a good concept to remember
 - otherwise, the connection atrophy and the concept can be eventually forgotten

Can you explain it?

Dendrits migration

- There are about 128^4 possible concepts – much more than c-neurons!
- Each dendrit has its **strength** (it is not the STW!)
- The strength is strengthen by a constant iff:
 - strength-gain-SVRule > 0
 - e.g.: Increase the value of strength by 3, if "susceptibility>0 and chem3>0"
- The strength is weaken by a constant iff:
 - strength-loss-SVRule > 0
- If the strength of a dendrit is decreased to 0 (or all dendrits comprising a concepts are 0, respectively) the dendrit may migrate to another input neuron



Outline

1. Gentle introduction to norms
2. Neural networks
3. **Possibilities and limitations**

Possibilities

Action sequences

- Action sequencing could appear theoretically:
 1. <hunger (high)> + <food> => walk left
 2. <hunger (high)> + <food> + <<it is near>> => eat it
- The second concept becomes valid after walk left is executed

Possibilities

Robustness

- Consider the following rule:
 - <hunger (high)> + <food> + <<it is near>> => eat it
- The rule is strengthen every time the norm eats
- When some food is rotten, a punishment is generated and the rule is weaken
- Despite the episode with the rotten food, after a while, the rule will fire again, because the STW exponentially increases to the LTW (while the LTW decreases only a little bit)

Possibilities

Generalization

- What is generalization?
 - *I find myself in the situation looking at a big green thing with staring eyes, which I've **never** seen before. Nevertheless, I remember from previous experiments that going up to things with staring eyes and kissing them is not good idea. So, I'll try something else this time.*
 - The creature can deal with situations, in which it has never been before!

Possibilities

Generalization

- What is generalization?
 - *I find myself in the situation looking at a big green thing with staring eyes, which I've **never** seen before. Nevertheless, I remember from previous experiments that going up to things with staring eyes and kissing them is not good idea. So, I'll try something else this time.*
 - The creature can deal with situations, in which it has never been before!
- But what about exceptions...?
 - *This horrible green creatures is my kind aunt, who gives me a lollipop from time to time.*

Possibilities

Generalization

- In this model, generalization is theoretically possible:
 - a) <norm> + <<it is opposite-sex>> + <sex-drive (high)> => run
 - b) <norm> + <<it is opposite-sex>> + <<it is my-parent>> + <sex-drive (high)> => stop
 - The rule (a) holds in both situations
 - 1. Consider a situation that the rule (a) applies for (B)
 - 2. When this happens, a punishment is generated and STW of the dendrit_{A=>run} is weaken and the dendrit_{B=>stop} is strengthen
 - 3. After some repetitions, it becomes:
 $LTW_{A=>run} + LTW_{B=>stop} < LTW_{B=>stop}$. Now, the rule (b) is firing in the situation (B) and STW of dendrit_{B=>stop} is strengthen as a consequence.
 - 4. Later, if the situation (B) occurs again, the rule (b) will apply.
 - 5. The rule (a) will still apply in the situation (A), because (b) does not fire in (A) at all.
- **Summary:** The dendrit from the more general concept to its desired action will have smaller LTW than the dendrit from the more specific concept to its desired action, and/or (perhaps) the more specific concept will inhibit the action of the more general concept. It means that the creatures will have learnt to discriminate between the two concepts.

Limitations

- The attention is focused only to 1 object
 - “subject-verb-object” is not possible
 - e.g. “put egg to hatchery”
- The norm can not deliberately switch attention
 - If a norm is more stimulated by a toy (which is attractive and could be moving) rather than by food (which is not moving), it will be fixated on the toy although it is dying of starvation
 - Can a connection between drive lobe and attention lobe help?
 - Must this wiring be hard-coded or can be learnt by reinforcement?
- The norm has no world-model memory
 - It has a procedural memory, but does not have episodic memory
 - It is not able to learn where the garden with some food is

Main limitations (reactive planning vs. Creatures' net)

	R.P.:	Creatures:
• Transition	SHRP + FSM	N. A.
• It behaves in the same way	probabilistic	OK?
• Rigid	-	N. A.
• Compromise action	-	YES
• Proscription	-	YES
• Modification of a behavior	BOD	-
• Concurrent behaviour	-	-
• Interleaving	-	-
• Sharp timeout	-	N. A.
• Adaptive	-	partially?
• Time-consuming design	-	KO!

End.

References

Creatures:

- Grand, S., Cliff, D., Malhotra, A.: Creatures: Artificial life autonomous software-agents for home entertainment. In: Johnson, W. L. (eds.): *Proceedings of the First International Conference on Autonomous Agents*. ACM press (1997) 22-29
- The Creatatures Developers Resource: <http://www.double.co.nz/creatures/index.htm>
...relevant information, but not updated!

Evolutionary computation:

1. J. Heitkoetter and D. Beasley: THE HITCH-HIKER'S GUIDE TO EVOLUTIONARY COMPUTATION: <http://www.faqs.org/faqs/ai-faq/genetic/>
2. Interactive Introduction to Genetic Algorithms: <http://ics.felk.cvut.cz/~xobitko/ga/>
3. John Koza's homepage: <http://www.genetic-programming.com/johnkoza.html>

Perception lobe

