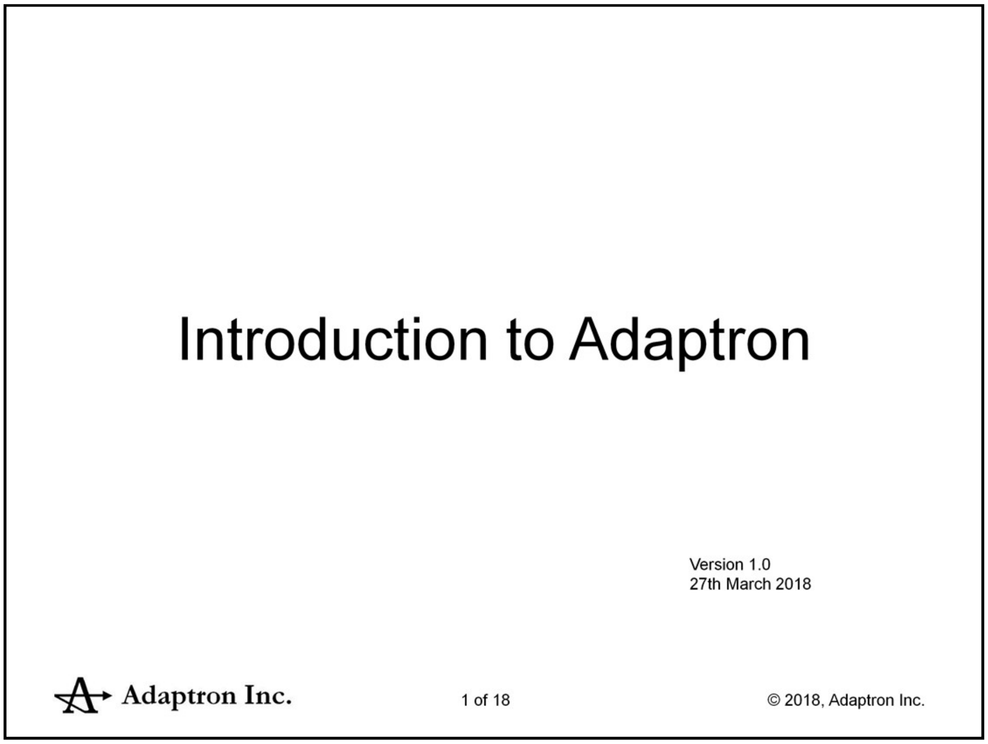

This introductory presentation provides a short overview of what Adaptron is and the objectives for its development.
It summarizes the key requirements and touches upon the subjects of:
- Autonomous robots and their motivations
- Stimuli and Responses
- Pattern recognition
- Binons and the artificial neural network
- Action control
- Short term memory to track what was perceived, what was done, where and when
- Thinking

Adaptron

- An Information processing model of learning and thinking
- A cognitive architecture for Artificial General Intelligence (AGI)
- General purpose robot control software

Adaptron Inc.

I originally described Adaptron as an information processing model of learning and thinking.

A more up-to-date description is a cognitive architecture for Artificial General Intelligence.

In more practical terms it is: general purpose robot control software.

Requirements

- Grounded on senses and action devices
  - Embodied
- Simple
  - Occam's Razor
- Deterministic
  - Explanatory behaviour
  - Completely observable

Adaptron Inc. 3 of 18

The Adaptron software is designed to operate in a robot body with sensors and action devices.

This means it is grounded on real world stimuli and behaviour.

One of the driving principles for its design is to be as simple as possible and still provide the functionality required.

Another requirement is that one can observe its internal functioning and explain in a deterministic way how it learns, acts and thinks.

This means that you will find all my definitions and descriptions are very strongly mechanistic / functional.

Requirements

- Compositional
  - Reuse simpler behaviour
- Learn quickly
  - Minimal repetition
- Biologically inspired

Adaptron Inc. 4 of 18

Compositional means Adaptron needs to build more complex behaviours out of simpler behaviours.

The simpler behaviours are learnt first and combined to achieve more complex behaviours.

Learning also needs to be based on as few experiences as possible.

Adaptron should not have to repeat an experience more than a few times to learn it, provided it is based on less complex behaviour that it already knows.

Biological evolution has already resulted in intelligent organisms, so I reuse any designs that nature has created that are useful.

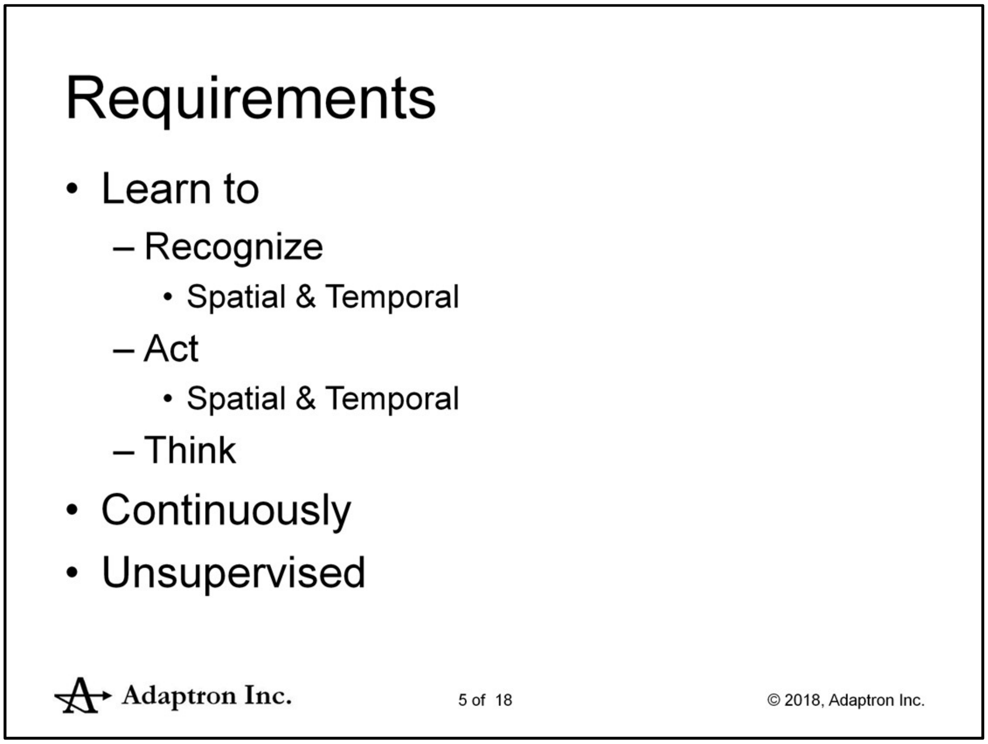

Adaptron needs to learn to recognize both spatial and temporal pattern.

It needs to learn to perform actions simultaneously on more than action device.

And it needs to learn to think.

It needs to operate continuously in the real world in an unsupervised manner.

By unsupervised I mean there is no training stage and then a test stage in its growth.

It must continuously learn the patterns in the world and adapt its actions based on its motivations.

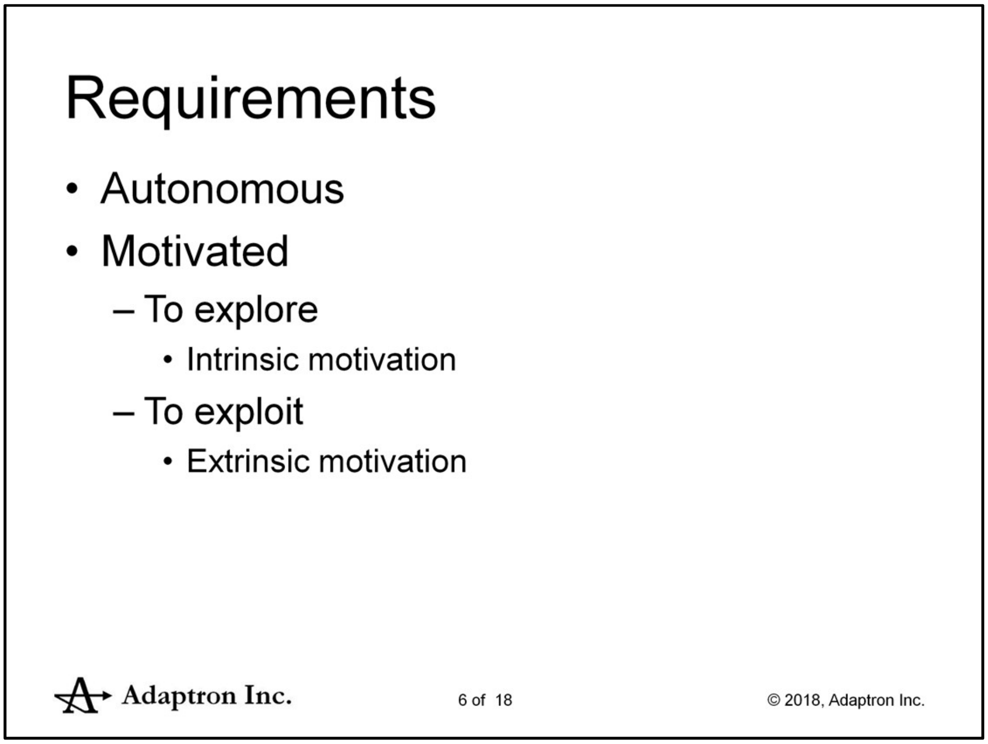

Adaptron must not be controlled like a puppet. It needs to act independently.

Its motivation to act must come from the need to explore and to pursue pleasant stimuli.

Exploration seeks out novelty and tries to avoid familiarity.

This is intrinsic motivation because it arises from inside the agent.

From a functional perspective it is just the comparison of current experiences with remembered ones.

If an experience matches a remembered one it is familiar, otherwise it is novel and worth re-experiencing until it becomes familiar.

Pursuit of pleasant stimuli and avoidance of unpleasant ones is called extrinsic motivation because the reward comes from external things.

However, which stimuli are pleasant and unpleasant is built-in when the software is programmed.

Types of Motivation

- Intrinsic [1]
  - Novelty
    - Interesting stimuli
  - Familiarity
    - Boring stimuli
  - Produces exploratory behaviour - curiosity
    - Orienting responses
  - Inherent in information processing
    - Wherever there is memory

Since motivation is such an important aspect of autonomous robot behaviour I cover this subject in more detail.

Intrinsic motivation results in exploratory behaviour such as play and pursuit of entertainment.

Surprising stimuli are novel. There is an interest in re-experiencing them.

Once they are fully explored and familiar they loose their interest and become boring.

Animals are naturally curious and at a young age most of their learning is based on intrinsic motivation.

The orienting response is mentioned an example because it is a biological reflex that attracts ones attention to novelty.

Unexpected changes in the environment are novel and they interrupt our behaviour.

Intrinsic motivation is an inherent property of information processing. New information is naturally compared to existing information which are memories.

Types of Motivation

- Extrinsic
  - – Pleasant
    - Rewarding stimuli
  - – Unpleasant
    - Punishing / painful stimuli
  - – Produces purposeful behaviour
    - Directed action
  - – Is built into a robot

Adaptron Inc.

Pleasant and unpleasant stimuli provide the extrinsic motivation to perform purposeful behaviour.

In humans they include the pleasure from certain foods, touch and other survival related motivations.

Extrinsic motivations would be decided on and built in to a robot by its creator.

However they could be left out.

In that case the result would be a robot that is always curious and continuously explores.

It would always be playing and seeking entertainment.

Stimuli

- Multiple senses
  - Multimodal / Independent senses
  - Multiple sensors – multiple properties
    - Independent
    - Dependent
  - Types of values
    - Magnitude
    - Symbolic

Adaptron Inc. 9 of 18

Adaptron must to be able to take in stimuli from multiple types of senses and sensors.

Examples of senses include sight, hearing and touch.

These three senses are independent of each other because they detect different sources of energy.

Within each sense there are multiple types of sensors and many of each type.

For example in sight there are two types of sensors. The rod sensors for brightness and the cone sensors for colour. And there are many sensors of these two types covering the surface of the retina.

Sensors can be independent such as the ones in your right ear are independent of the ones in your left ear.

But if two sensors of the same type are beside each other, they are dependent, they have an adjacency relationship.

Biological sensors measure the magnitude or intensity level of the property they detect.

But a robot could have more intelligent sensors that identify things and return symbolic information such as car, tree, house, pedestrian, bicycle, post, stop sign, etc.

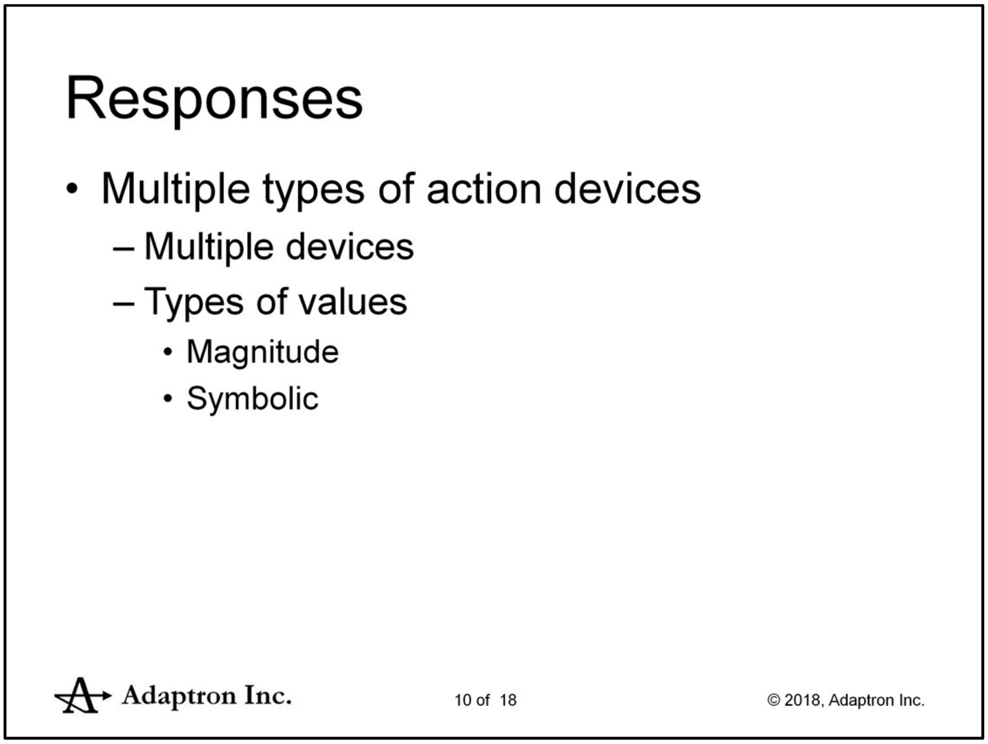

Adaptron needs to have an effect on its environment by producing responses on its action devices.

There are many types of such devices.

Biological action devices are usually muscles that might control limbs, eyes or the larynx for ones voice.

Robotic examples are motors to control limbs or wheels, lights and speakers.

And a robot could have many of each type.

Simpler action devices are given magnitude responses indicating the change they should make.

More intelligent action devices are given symbolic responses such as desired position, a final result.

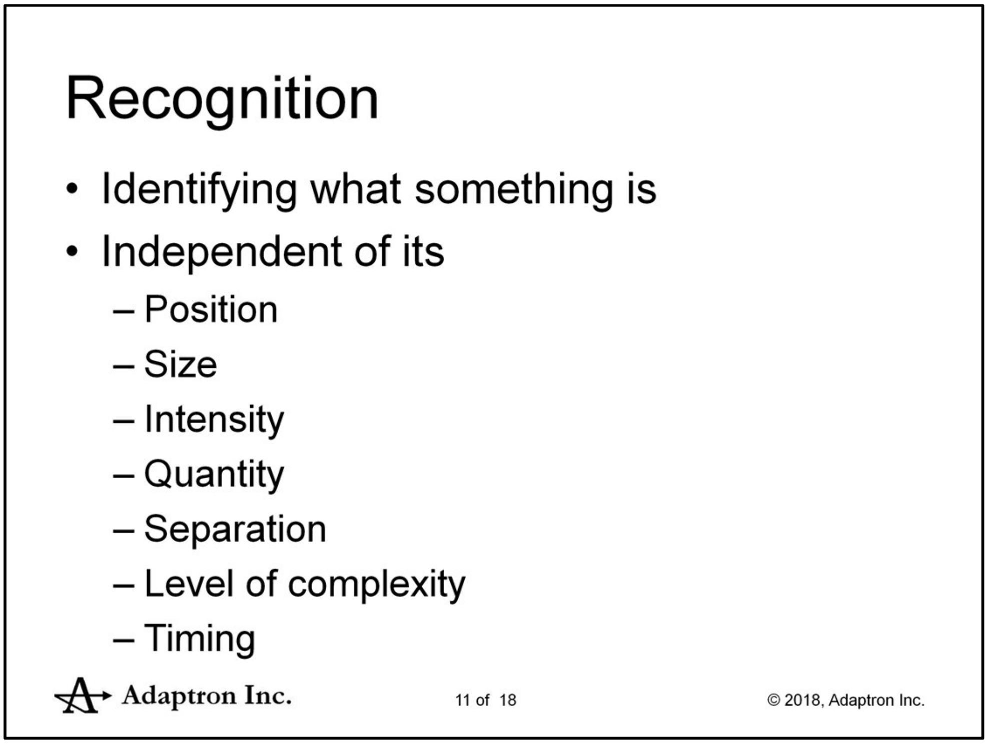

Recognition is the process of identifying what something is.

And Adaptron needs to learn to recognize things in its environment.

And once such things have been learnt it must be able to re-identify them no matter where they are, how big they are, how bright they may be, how many there are, how far apart they are or how far away they are.

It must be able to do this for things at all levels of complexity.

It must also be able to recognize things as they change in time. A rotating object is a good example. It might look quite different from one angle to another.

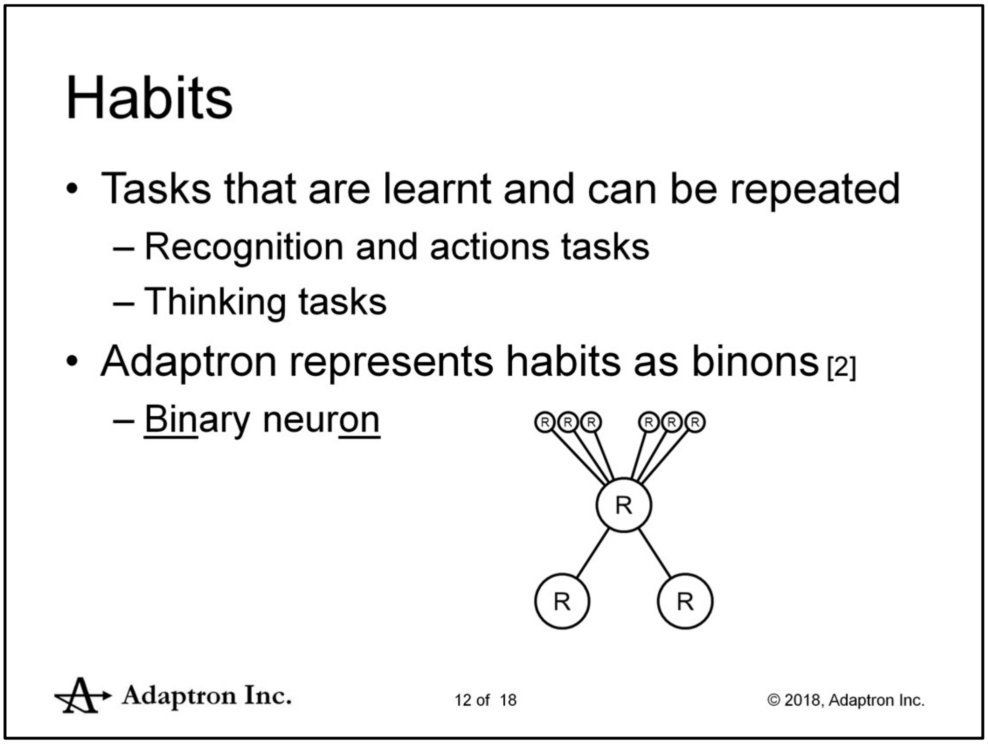

Adaptron learns and repeats habits. Habits are tasks that can be performed. These include habits to recognize things, to do things and to perform thinking.

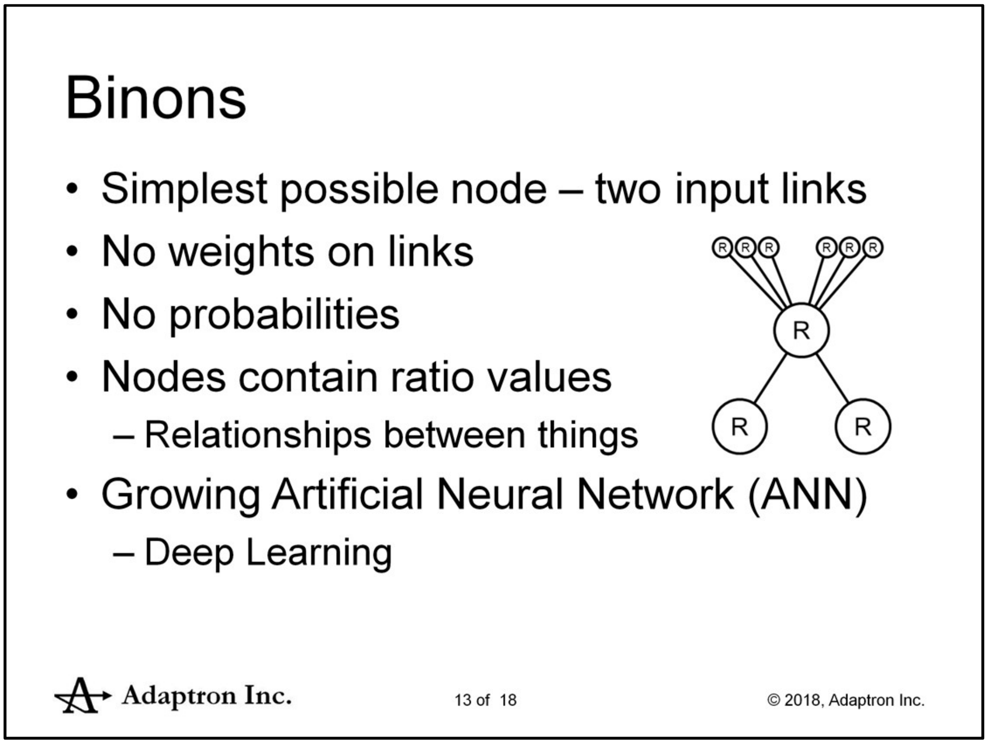

Binons are the simplest possible artificial neural node that can combine information and represent things that can be done.

And they represent things that are recognized and actions that are performed.

They are combined in a growing multi-level tree structure which reuses lower level binons to form more complex binons.

This is a deep learning architecture.

It differs from many such deep learning architectures in that it has no weights on the links and makes no use of probabilities.

This allows one to explain exactly how it makes decisions and behaves (i.e. deterministic)

However the binons contain ratio values that quantify the relationships between things.

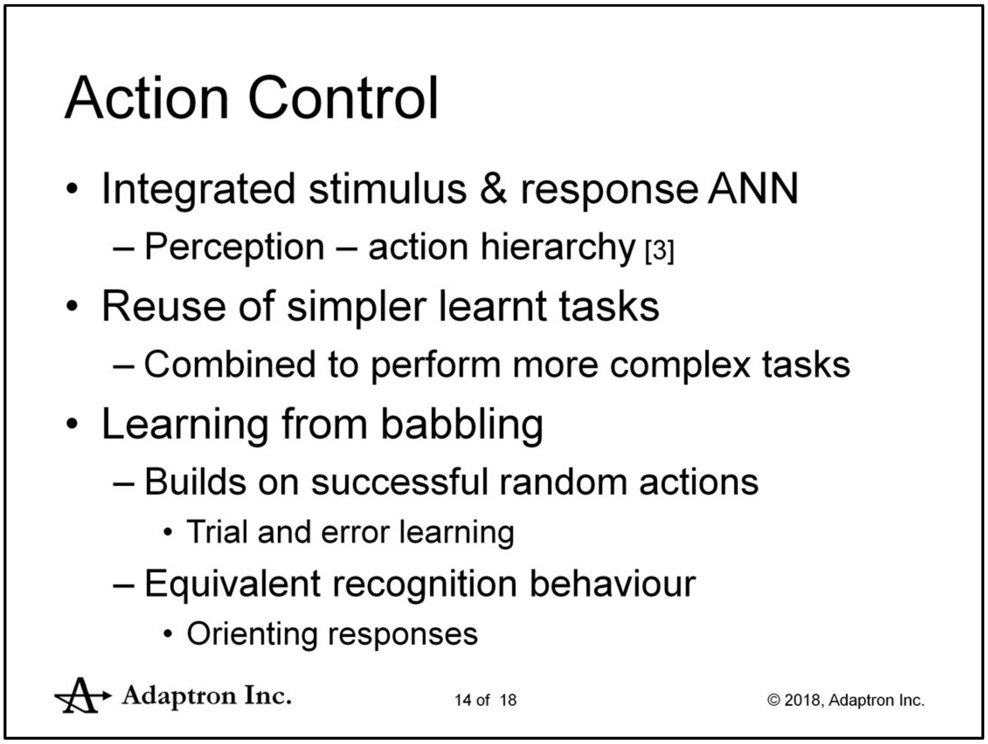

Adaptron' s ANN is a tightly integrated network of the things it recognizes and the actions it performs.

Actions are represented as binons just like things that are recognized.

The same rules and processes used for learning to recognize things are also used for learning how to perform more complex action from simpler ones.

The one important difference is that action binons result in responses while recognition binons result from stimuli.

An example of the similarity between learning to recognize something and to perform an action is babbling.

Babbling is the random performance of actions and the memorization of those that produce interesting results.

These are then repeated to obtain the same results or combined with other learnt actions.

Most people are familiar with speech babbling. It is a subset of motor babbling.

The orienting response is the recognition equivalent to motor babbling.

It is a biological reflex that attracts attention to unexpected changes in stimuli.

Initially it is quite random because all stimuli are unexpected. One has no experience, no memory.

The orienting response causes a change in attention which is then learnt and repeated to obtain the same stimuli.

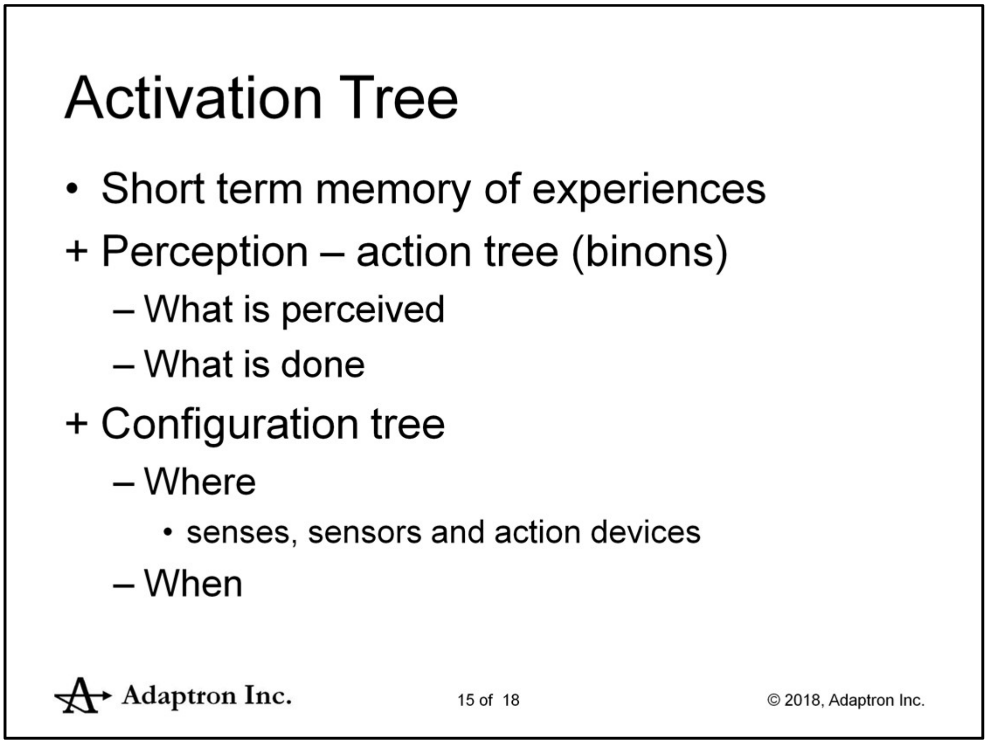

Adaptron maintains an activation tree that is the short term memory of its experiences.

The activation tree combines the What, Where and When information of an experience.

What something is and what is done are represented as binons in the integrated perception – action tree.

However where and when something is done also needs to be known. This information is in the configuration tree.

It is a binary structure that identifies on which senses and sensors things are perceived and

on which action devices things are done.

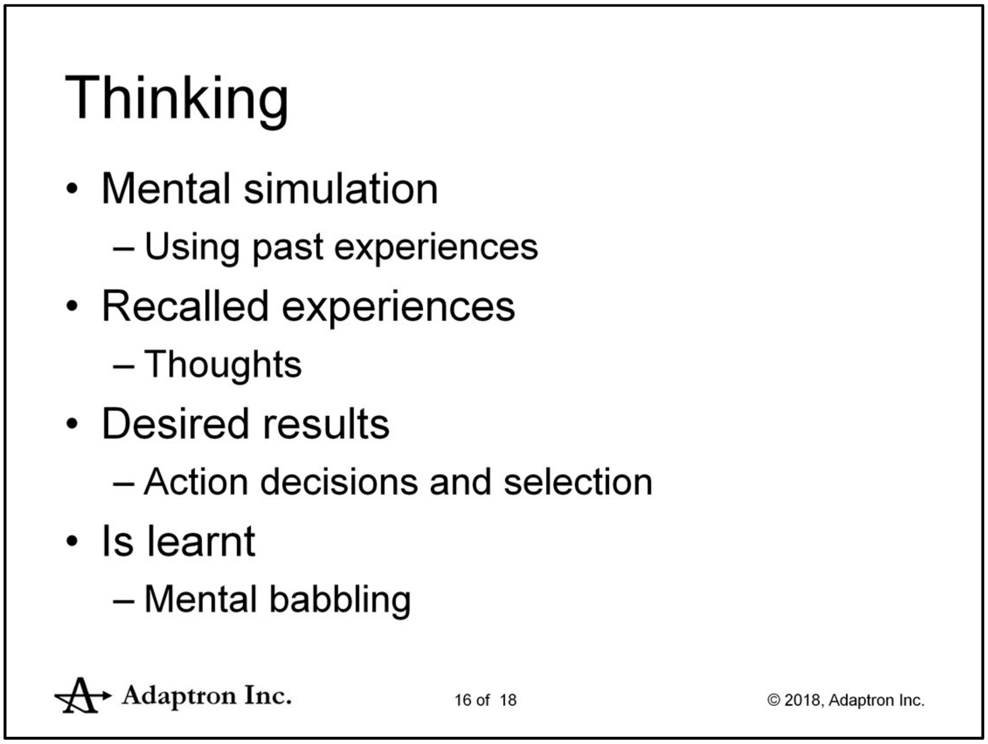

Adaptron needs to think. It needs to use its experiences as its model of reality.

It uses this model to mentally simulate the possible consequences of its actions before doing them.

But rather than pre-program this thinking process Adaptron must learn it.

The hierarchical ANN of binons representing ones experiences is the environment for thinking.

It is the source of thinking’s stimuli (ideas or thoughts). And thinking’s actions are the focusing of attention on these binons to recall them.

These actions result in the traversal of the network. Up the tree for more specific things, down the tree for more general concepts and across the tree for associated things in space or time, i.e. what is expected to happen next.

But thoughts are not just recalled experiences. They also include stimuli that represent how well the thinking process is working.

They include information such as “I know I knew it but I can’t remember it”, “It’s on the tip of my tongue” and “I know I don’t know that”.

So a train of thoughts comes to an end when an action is thought about that would result in something desirable. The decision is then made to perform that action.

Thinking starts with mental babbling, just like recognition and action learning.

Detailed Presentations

- Senses, Sensors and Action Devices
- Perception
- Recognition
- Binons
- Hierarchies
- Learning
- Experiences
- Temporal Recognition
- Action Learning
- Thinking

Adaptron Inc. 17 of 18

This is a summary of the content of the more detailed presentations as currently planned.

The first 3 presentations provide mechanistic and functional descriptions of their subjects to set up the foundation for the architecture and design of Adaptron.

Presentations 4 through 10 cover binons, how they are structured in the ANN and how they work to perform spatial and temporal pattern recognition and action learning.

The subject of thinking and how binons can be used for that purpose is described in the last presentation in more theoretical terms.

I have not had time to refine the thinking cognitive architecture and prove it in software.

References

[1] Baldassarre, Gianluca and Mirolli, Marco Editors. (1998). Intrinsically Motivated Learning in Natural and Artificial Systems, Springer, Publishers

[2] Martensen, B. N. (2013). Perceptra: A New Approach to Pattern Classification Using a Growing Network of Binary Neurons (Binons). In R. West & T. Stewart (eds.), Proceedings of the 12th International Conference on Cognitive Modeling, Ottawa: Carleton University.

[3] NIPS 2007 Workshop. Hierarchical Organization of Behavior: Computational, Psychological and Neural Perspectives, http://www.princeton.edu/~yael/NIPS_workshop

Adaptron Inc.

1/ A good book about intrinsic motivation

2/ My paper on binons and

3/ a link to excellent presentations on the Hierarchical Organization of Behavior.