

Activity 2

Brain-in-a-bag: creating an artificial brain

Age group successfully used with:

8 – adult

Abilities assumed:

answering general questions,

Time:

20-30 minutes as lecture format,
1 hour as workshop

Size of group:

anything from 7 to hundreds

Focus

What is a neural network?

How can a working artificial neural network be created?

Summary

After explaining how neurons in the human brain work we create an artificial version that is programmed to play Red-black snap. We then test it against other artificial brains.

Aims

This activity explores how biology and specifically animal brains can be the inspiration for a new way to program computers and so give them intelligent abilities. In particular it aims to give an understanding of how neurons work and so how this architecture may be copied to create an artificial neural network that can lead to human-like behaviour.

Technical Terms

Neural networks, biologically-inspired computing, artificial intelligence.

Materials

- Powerpoint with slides showing a real neuron.
- A single pack of cards (ideally jumbo sized) or alternatively normal black and red card could be used.
- A pair of 'position' sheets showing positions for the cards.
-

A full class activity could involve making 4 brains using 28 people with others holding and turning the playing cards. Lecture style would involve building a single brain with 7 volunteers. For each brain (requiring 7 people to operate) you need:

- 6 tubes
- 6 lengths of rope around 1 or 2 m long
- 1 set of 7 neuron rules cards

- A bag holding the above

What to do

Preparation: Before the session put in each brain bag the six tubes and rope. To make things easier when creating the brain, tie 3 of the lengths of rope into a Y shape, knotting them together in the middle of the Y. Repeat this with the other 3 lengths of rope.

The Grab: Explain that you are going to create a working artificial brain using your brain-in-a-bag kit. Wave one of the bags of materials around.

Explaining neurons:

To make a brain the class first need to understand how a real brain works. Show the slide of a brain cell. The following is a simple description. It could be linked directly to a longer biology session where neurons are described in more detail but that is not necessary for the purposes of this session.

- Biological brains are made of *neurons*. These are just brain cells. The human brain contains billions of neurons all connected together.
- Each neuron has a central core and then lots of branches from this central core. The branches connect neurons together. Each neuron will be linked to lots of others.
- Neurons communicate with each other by sending simple messages down the branches. These messages are just pulses of electrically charged chemicals.
- Each neuron operates by following a simple rule. The neurons that connect to it send it messages. When enough messages have arrived (and so a threshold is reached) the neuron fires. What does it mean for a neuron to fire? Simply that it sends a message to all the neurons it connects out to. Each neuron has its own threshold that determines when it fires.

[The model of a balloon being blown up as it receives messages until it eventually pops (the neuron fires) as shown in the slide is one way to explain this. You could even try it with a real balloon and pump.]

That is enough biology. The point is that ultimately everything that happens in your brain boils down to billions of neurons doing that kind of thing. Doing homework, watching TV, falling in love is all neurons sending chemical messages to each other, over and over again.

The Set-up:

Thanks to neurobiologists and cognitive scientists, we understand more and more about how the brain works. As a result computer scientists can now try to build artificial brains that work in a similar way.

Explain that you are going to create a brain using your brain building kit. Each brain will be able to play Red-Black Snap. If doing this in a lecture format, one brain and

seven volunteers is sufficient. For a whole-class activity you can make several brains that then play snap against each other.

Get your group of 7 volunteers out to the front and hold up a brain-in-a-bag and explain that this is all you need: rope and toilet roll.

The first step is to set up the physical structure of the artificial brain. Your seven volunteers for each brain will act as neurons and provide the energy. The ropes act as the branching structures that connect the neurons together (the Axons). Take one set of Y-connected rope and get three of the volunteers to hold the ends. A fourth person holds the knot in the middle.

The person holding the base of the first Y also holds a second piece of rope – the base of the second Y. Two new people hold the other ends and the final person holds the knot of this second Y. You should end up with the layout of Figure 1.

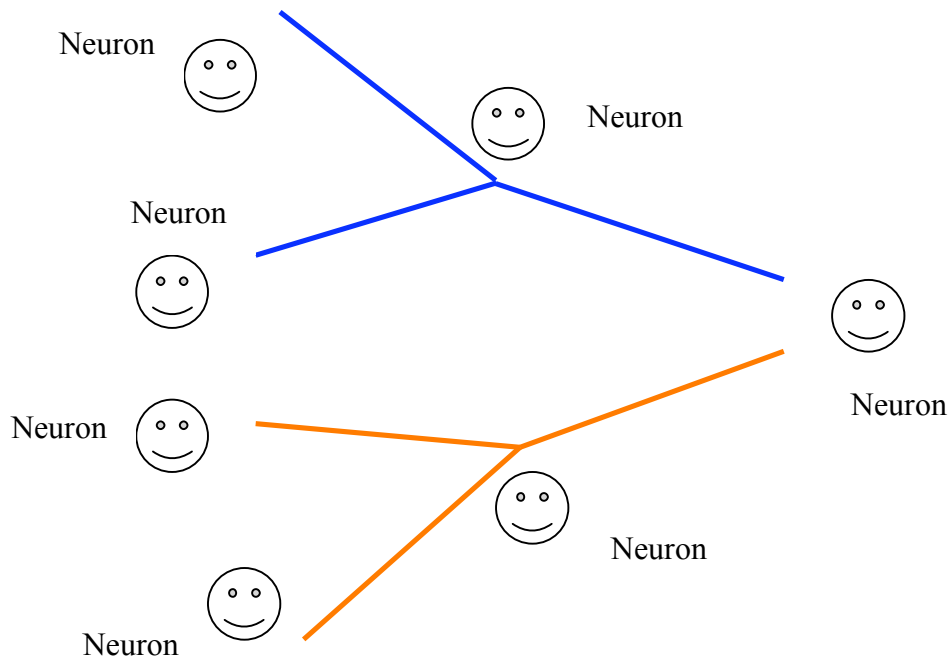


Figure 1: Laying out the rope and positions of the volunteers

You now have a brain but it can't do anything as the neurons have no way to communicate. Explain that the next step is to give the neurons chemical messengers. That is what the tubes are for. Each of the four people at the tops of the Ys need a tube threaded on their rope. Each of the two people holding the knots also need tubes. Thread them up to the knot from the person at the base of the Y.




Figure 2: The initial positions of the tubes

Explain that those with tubes can send messages to other neurons by firing the tube down the rope (the rope needs to be pulled tight). Explaining this may trigger some random brain activity in your newly created brain as tubes go flying about!

You have the neurons, an energy source (the volunteers) and a way of passing messages but still the brain won't do anything useful – even if there is signs of random brain activity! It can only work in an organised way if the neurons have rules to follow. The rules would normally be learned over a lifetime. You will speed things up by providing pre-programmed rules.

Give each person a Neuron rule card to follow (see Figure 3). The four people at the end of the Ys (numbered 1 to 4) are eye neurons. They are connected to the eyes and they are the only way that the brain can sense the outside world. Their rules are about “firing” if they see a particular colour in a particular place (more on that below). *The important thing is that the rules about the same colour are on the same Y of the brain.*

By “firing” we just mean sending the tube down the rope to the connected neuron.

The next two neurons (5 and 6) are deep in the brain. They are not connected to the outside world, only to other neurons. They have a rule to fire once their threshold is reached. That is, they send their tube onwards *only* when have received TWO tubes from connecting neurons.

The final neuron is a snap neuron (the kind everyone developed around the age of 4!). It is connected to the mouth of our artificial creature and is the only way the brain has to communicate with the outside world. Its threshold is a single tube arriving. When it fires, the mouth shouts Snap as loudly as possible.

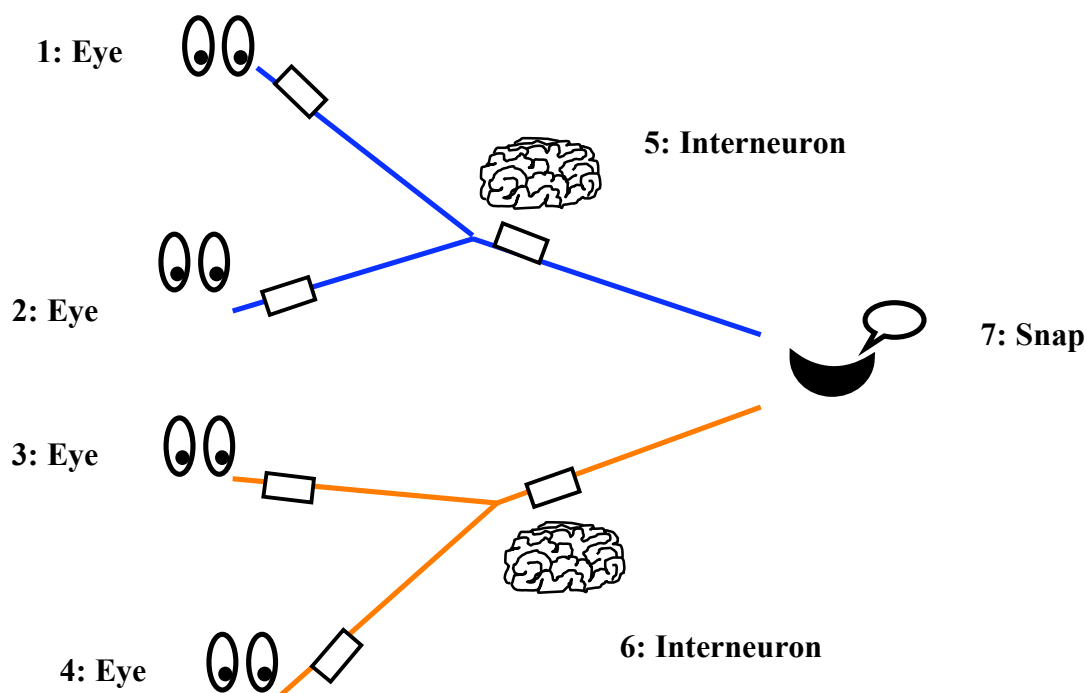


Figure 3: Placing the rule cards

Finally, you need something for the brain to look at. Either photocopy the two sheets provided onto a single sheet with two positions marked on, or get two more volunteers to be the two positions. The first is the ‘position 1’ referred to in the rules, the other is ‘position 2’ of the rules. You will place red or black cards in these positions.

Explain that the eyes have to watch only the position their rule tells them to watch. They should also just look for the specific colour instructed. The other people in the brain are neurons with no ability to sense the world. They should therefore turn their backs to the cards. They only react to tubes arriving and their threshold being met, not to what is happening out in the world.

Testing the brain:

The artificial neural network is now ready to be tested. Place a red card in position 1. Say out loud what you are doing to help the volunteers get it right: “A red is being placed in position 1”. Neuron 1 should fire but nothing else should happen. Now place a red card in position 2. Neuron 2 should fire. This should cause neuron 5 to fire, which should in turn cause neuron 7 to shout Snap! The brain may need a little prompting to work properly. If so you can make jokes about that is why slow brains need coffee to get them going. If anyone drops a tube or the end of a rope you can make a joke about brain surgery being needed.

At this point explain that after neurons have done a round of firing there is a resetting period as the chemicals return. The tubes all go back to the start.

Repeat this with two black cards. The other side of the brain should now fire. The brain could of course be just shouting “Snap” whatever happens. To check this next put up a red and then a black card Finally, put a black card in position 1 and a red in position 2. You should now have shown that the brain is fully working.

Playing Snap:

You can now try it for real. Make sure only the eye neurons are looking. Quietly turn cards over at random from the pack and see what happens.

If you are building several brains in parallel then at this point you can actually have a proper game of Snap between them. Giving points for the brain that shouts snap first. Deduct points for any spurious snaps or other nefarious activity that happens.

Summing up:

Explain that you have created an artificial brain together that works in a similar way to a real brain. When you play snap a similar chain of events goes on in your brain. Here we did it with physical things (rope and tubes), but it could just as easily be simulated in software with virtual neurons sending electronic messages to each other. That is essentially all a neural network is, though like real brains, real neural networks have a way to learn rules for themselves.

We only had 7 neurons. With more we could have created a brain that plays snap properly not just red-black snap. Imagine what an artificial brain with a billion neurons could do. Computer Scientists are working on it!

Variations and Extensions

A good, short follow on activity is to watch the affective robot video. The robot in the video has a software neural network brain so works in essentially the same way our brain-in-a-bag worked.

A further longer activity is to get groups to create their own neural networks out of rope and tubes to recognize animals – looking for things like stripes, claws, etc

Ropes and tubes can also be used in a similar way to illustrate logic gates and program flow of control.

Further Reading

Computer Science and the cell

Lots of articles on the links between computer science and cell biology

<http://www.cs4fn.org/biology/cells/>

Can computers help understand the brain?

How might you do statistics using neural networks?

<http://www.cs4fn.org/biology/basalganglia.php>

Life in the Undergrowth: Bee Faces

How complex do neurons need to be to recognize faces?

<http://www.cs4fn.org/biology/beeface.php>

The Centre of the Cell

Interactive activities linked to the national curriculum on cell biology

<http://www.centreofthecell.org/>

Links to other activities

The affective robot video

Watch a video of a robot with a neural network brain.

This is a good way to introduce both affective programming, and show what can be done with a small neural network.

Spot the Dalmation with a neural network (forthcoming)

Design your own neural network to recognize animals.

This is a good way to build deeper understanding of neural networks after the 'brain-in-a-bag' activity.

The intelligent piece of paper

Take part in a test of intelligence against an intelligent piece of paper!

This is a good introduction to what a computer program is, and also to start a discussion on what it would mean for a computer to be intelligent.

Where is the computer (forthcoming)

Spot all the computers in the room including the ones being worn.

This can be used as an introduction to the intelligent paper activity.

Sweets learning computers (forthcoming)

Make a computer that teaches itself to play a game perfectly.

A computer can be programmed to work the rules of how to play the game out for itself (so where is the intelligence then?) This activity is one way to illustrate how that can be done.



Conversations with computers— *The Turing test* (CS Unplugged activity)

How would you tell if a computer was intelligent?

Having possibly dismissed the idea that something following rules like the 'intelligent' paper is intelligent move on to explore how you would tell if a computer was intelligent or not.

Neuron 1



An 'eye' neuron



**Fire on RED card in
Position 1 only**

Neuron 2



An 'eye' neuron



**Fire on RED card in
Position 2 only**

Neuron 3



An 'eye' neuron



**Fire on BLACK card in
Position 1 only**

Neuron 4



An 'eye' neuron



**Fire on BLACK card in
Position 2 only**

Neuron 5


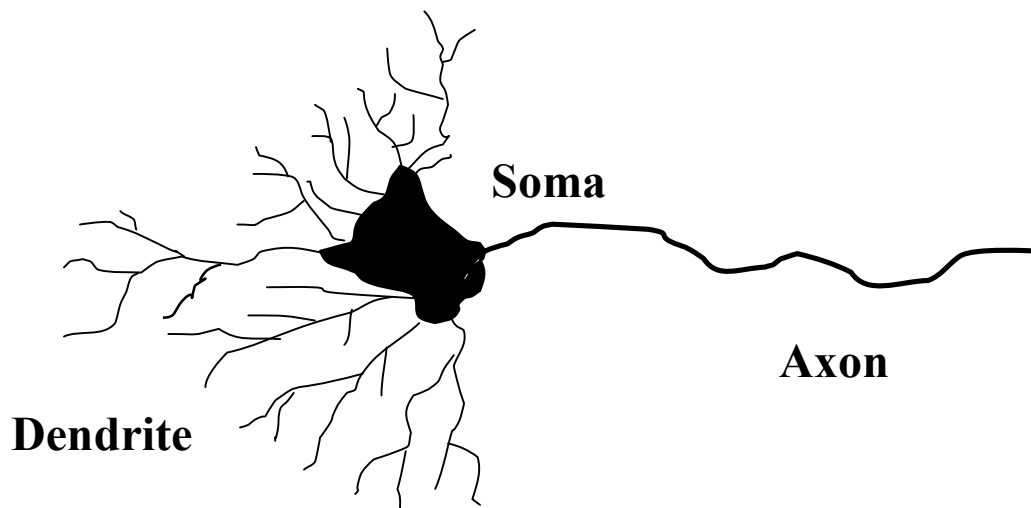
The interneuron



**Fire on 2 pulses coming
in only**

Neuron 6

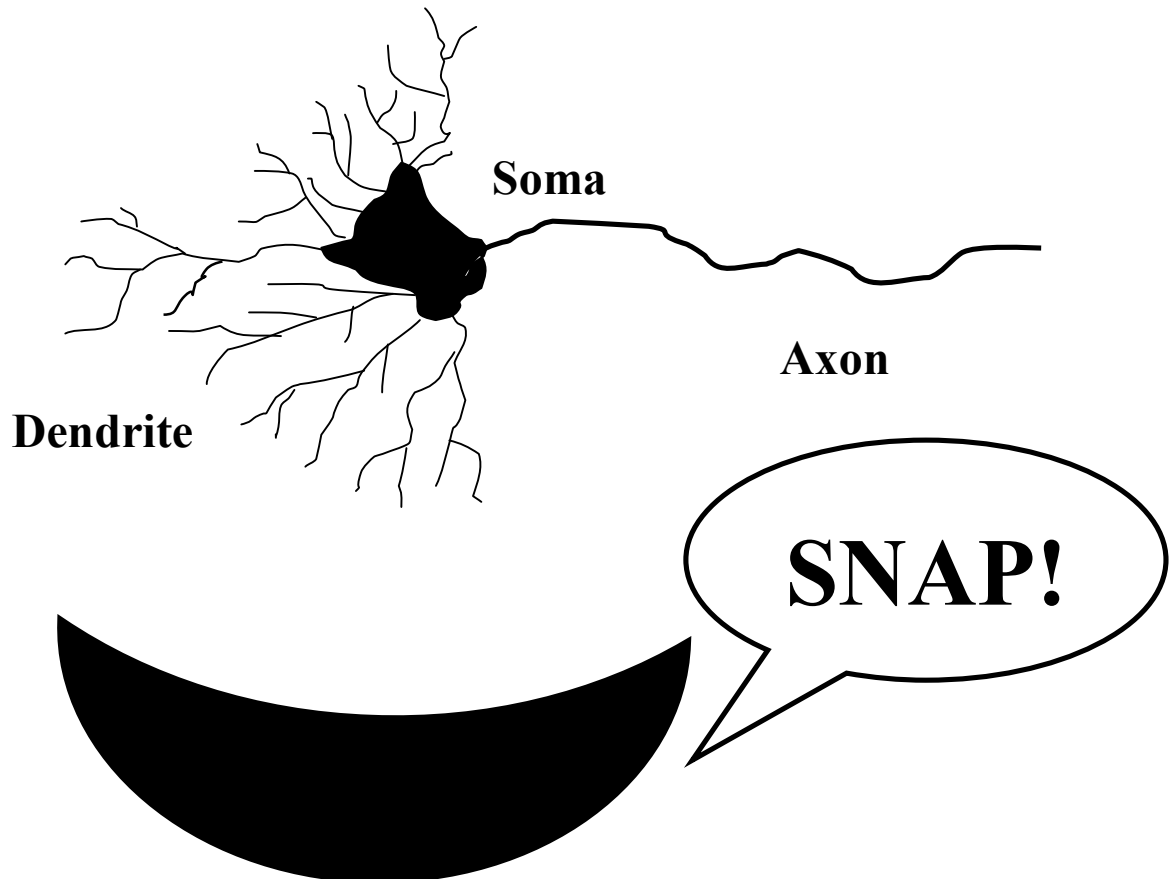
The interneuron



**Fire on 2 pulses coming
in only**

Neuron 7

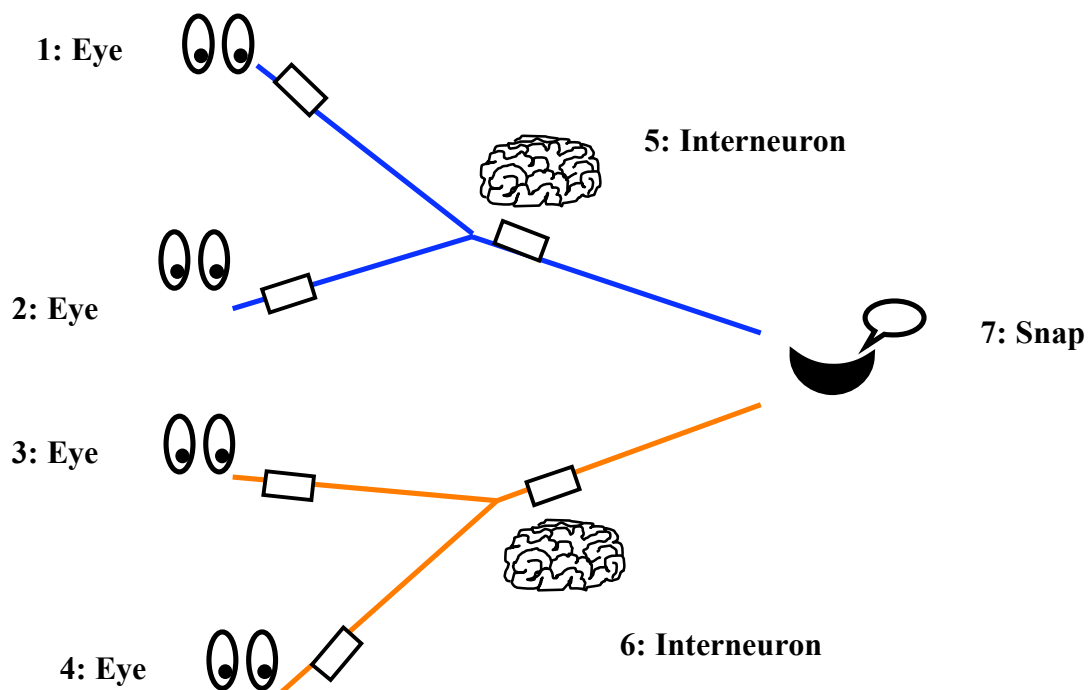
The 'snap neuron'



Fire on at least 1 pulse
coming in and then

Shout **SNAP!**

Wiring diagram for neurons 1-7



Position 1

Position 2

