

Cognition Ramblings

CBPrice 03-02-23 Overview without references.

Introduction

When we exist in our real world and move around and interact with this world, we can think of three sorts of things going on in our minds. First, we *perceive* the world (through vision, hearing, touch etc), then comes *cognition* where we make sense of what we perceive and perhaps decide what to do next, and finally there is *action* where we enact what we have decided to do. Existing in a game world is quite similar, once we become *immersed* in the game world where at this point, we are no longer aware of the mouse, keyboard and screen.

This whole process of existent depends heavily on our perceptions, and the most important of these is perhaps vision, since visual processing occupies a dedicated channel in our cortex which has a privileged location, Fig.1. Over the last hundred years, perception (especially the human visual system) has been widely researched and understood, so has our action systems (leg and arm muscles, speech). However, the *cognition* component, how our minds make sense of our perceptions and decide what to do next, is still the subject of theoretical and experimental research.

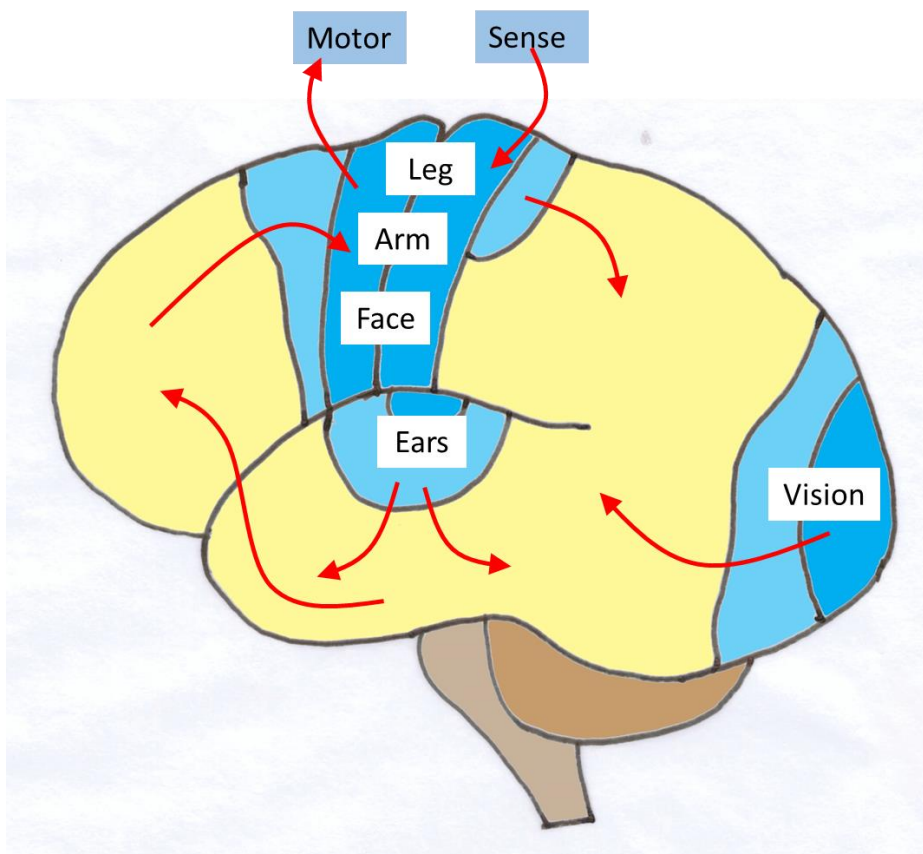


Figure 1. Shows the anatomy of the human brain. Blue areas are associated with sensory input and output, yellow areas are where various inputs are 'associated'. Red arrows show information flow. The top blue areas are concerned with paired sensing and motor action for various body locations. You can see there are lots of inputs and motoric outputs, but huge association areas which is the locus of thought.

Classical Cognition

So how do we make sense of the world? One could think the study of cognition is informed by psychology, but this is not the whole story, in fact classical cognition is greatly informed by computing, the work by Newell and Simon in the 1960s on the "General Problem Solver". Of course, as soon as the computer enters the discussion, concepts such as *processing*, *representation*. Enter our minds. The word 'processing' needs no elaboration, but the word

'representation' needs some thought. Let's step out of our brain and think of a computer and go to the fundamental component of a computer – wires! Look at Fig.2, on the left is a bus made of just 2 wires; each wire can be on (1) or off (0), so the pair of wires can be in the four states shown. That's all a computer 'knows' and these wires carry these four states into electronic components like adders. Now here comes, what I believe to be, some real magic. Look at Fig.2 on the right. Here you see the same 2 wires which the computer 'knows' but all stuff highlighted in yellow is *in the human mind*.




Figure 2. Left shows two wires of a computer bus with the 4 possible states of wires on or off. On the right are some human **representations** of these states highlighted in yellow. The computer is blissfully unaware of these.

So, the first *representation* is as denary numbers; if the left column of binary is taken to *represent* 2's and the right column is taken to represent 1's, then together these four states *represent* the numbers 0 to 3. The computer has no knowledge or understanding of 0,1,2,3 since it only knows (0,0), (0,1) ... (1,1). But we could choose to use these four states like the letters shown or even whole words, that's for humans to decide.

Let's now turn to **classical cognition**, where did it come from and what is it? Well, computer scientists and biologists were talking to each other. The neuroscientist McCulloch worked with the computer scientist Pitts to create the first mathematical model of the brain cell, the neuron. Then computing founding fathers Rumelhart and McClelland conceived of massively parallel neural networks with many connexions between neurons. Human thought was conceived as interactions between these neurons, where one neuron would excite or inhibit other neurons. Knowledge was conceived as residing in connexions between the neurons distributed over vast areas of the connected network.

But back to Newell and Simon's 'General Problem Solver'. They used computer representations which they called *symbols* which were associated with physical (world) *signals* (think perceptions). They combined the symbols into structures which they then manipulated, e.g., using logic.

Then we come to the work of Anderson who developed a model of memory and cognitive processing called 'Spreading Activation'. You can think of this as an overlap between the neural network concept, and the concept of symbols representing some real-world experience (perception). Each node in the network is a symbol, and the connexions between the nodes can grow or shrink according to how close the symbols are related. Let's take an example, you are out in the countryside, and you say to your mate "Look, the eagle is in the nest. According to classical cognition, what happens in your mate's mind looks like Fig.3. There are a load of symbols, drawn as circles and assume these all start off with some weak connexions (these are not drawn, I did not have patience). Hearing the word 'eagle' triggers the associated symbol and then this sends out signals to other related symbols. This process is spreading activation. Finally, the network settles down to some equilibrium state at which point your mate understands the words you said.

You may find this explanation attractive, and it is since it can easily be programmed (that's no surprise since the model's origins lie in computing), but it has some serious problems.

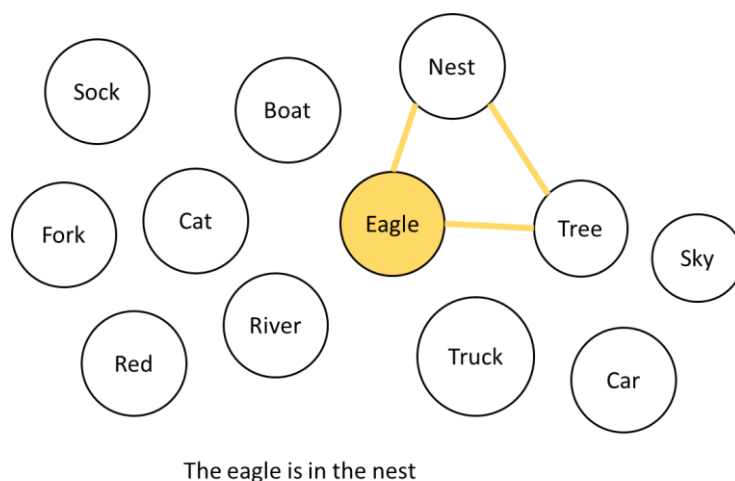


Figure 3. Spreading activation network triggered by hearing the word “eagle” and activation spreading to “nest” and “tree”. All circles are **symbols**.

Consider an example where these symbols do not work. Bob needs to change a broken light-bulb; imagine you say this to your mate “Bob stood on the chair to change the ceiling light bulb”, no problem. But if you said “Bob stood on the can-opener to change the ceiling light bulb” your mate would be very worried. But the spreading activation network would work for both cases, even though one is meaningless.

The problem is easy to see; the chair has the **affordance** of providing enough increased height to reach the ceiling, whereas, in this situation the can-opener has the **affordance** of inducing pain. The symbol ‘can-opener’ in the activation network is an **abstract** symbol and contains no information about the height of the object.

1. Bob stood on the chair to change the ceiling light bulb
2. Bob stood on the can-opener to change the ceiling light bulb




Figure 4. Example of where the classical symbol approach fails.

The problem is that these symbols are *amodal*, they contain no information about the real-world object affordances. Also, they are *arbitrary*, we are not told how the symbols actually emerge in the mind, it seems they are chosen by the researcher working on a particular problem. Finally they are *abstract* and bear no relationship to the real world.

Grounded (Embodied) Cognition

This is a contemporary theory of cognition associated with people like Gibson, Glenberg, Barsalou, Zwann and others. So, what’s new? You may have realized in the classical approach that there was a huge separation between perceptions in the real world and what was going on in the mind. Rather like mind-body duality where the body provides sensory

input *in one direction* to the mind, then the mind takes over (superior to the body) and does the real business of thinking.

Embodied cognition seriously challenges this model and proposes that the mind and body are intimately connected. There is no input-processing-output, but mind and body influence each other rather like in a *feedback loop*. This does make sense; we perceive something in the real world, like a door which we need to open, we know the handle is what we need to grab, we grab it and the door opens and we perceive the open door. Let's return to 'symbols' and see how to improve the classical model. This is sketched in Fig.5




Figure 5. Embodied representations of object contain various *modes*; images, sounds and also words.

You can see how each 'symbol' is no longer abstract, it comprises a bag full of images, sounds, and even words, all of which have been laid down in memory by our bodily experience of living in the real world. There is nothing *abstract* about these symbols, they have real *modal* content, and since they come from the real world, they are certainly not arbitrary, something we have made up. But there's more to embodied cognition than this, and what comes next could be very useful for a theory of game level design.

Affordances and Affordance Meshing

Affordances

We have seen the concept of *affordances* in the Bob the Builder example above, where the chair had the affordance of 'giving height'. Objects have many affordances (and I will suggest an approach to establish a classification below), but sticking with a chair, for humans it has the affordance of 'being sat upon' but not of 'being eaten'. A mouse does not engage with the affordance of 'being sat upon', for that little critter the chair has the affordance of 'can be run across'. The concept of affordances was introduced by Gibson, as that property which captures the relationship between perceived and actual objects, and also the way the body interacts with those objects. I hope you see the relevance to game-level design emerging.

Affordance Meshing

Returning to the two Bob the Builder sentences, one made sense and the other didn't. Let's see why; it's all to do with the affordances of objects in the situation, and how these are related (or not). The process of bringing affordances into relation with each other is called *affordance meshing*. Here's the affordances of the objects in play

Bob	
is 167 cm high	
can walk	can be punched
can climb	
can stand on	can be stood on
can talk	can be listened to
friendly	
skilled	

Chair	
is 50 cm high	
	can be stood on
	can be sat on
	can be moved
	can be picked up
furniture	
skilled	

Can Opener	
is 1.5 cm high	
can open cans	can be stood on
can inflict pain	can be sat on
	can be moved
	can be picked up
useful tool	

Table 1. Affordances for some objects in the Bob the Builder example.

So, for the sentence “Bob stood on the chair to change the ceiling light bulb” we have the following meshes:

{Bob: can walk, can climb, can stand on, is 167 cm high + Chair: can be stood on, can be moved, is 50cm high}. That meshed with the fact that the ceiling is over 200 cm high gives the sentence meaning.

Affordance Meshing applied to Images

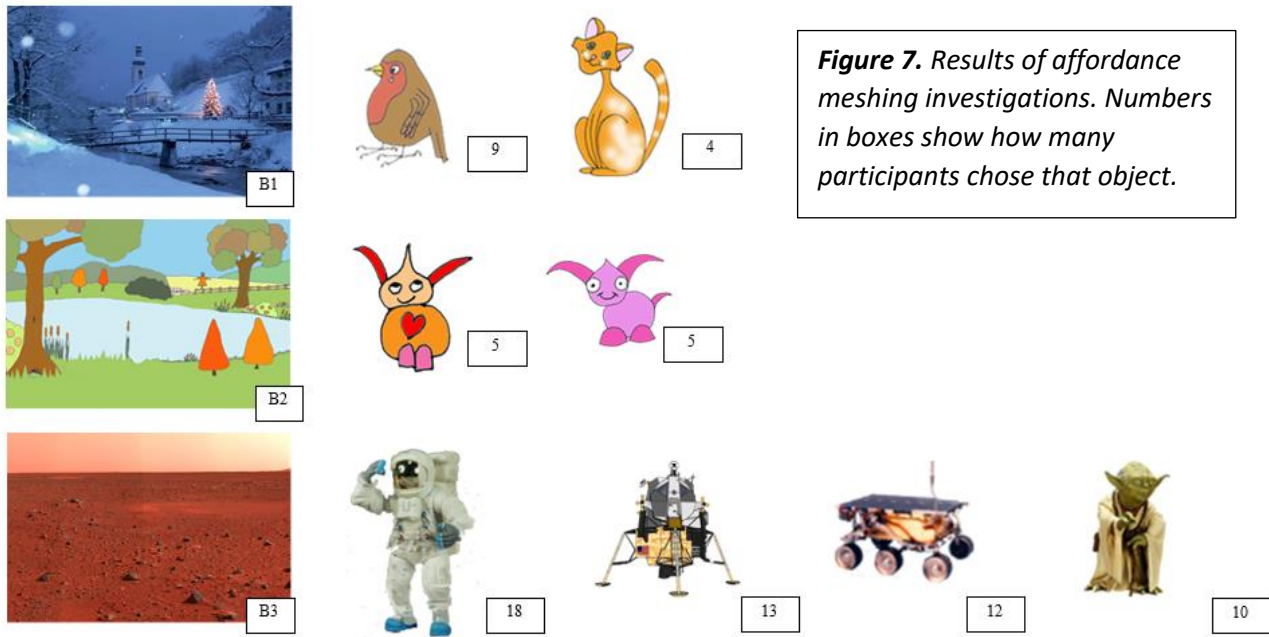
The above discussion has focused on language, because that was the original context for this theory to be developed, especially for language *comprehension*. Recently I have extended the theory to language *composition*, especially to writing stories using images. Now you could contribute to extending this theory to computer game level design. Let’s have a look at some of my research. The first investigation asked viewers to answer the question “What happens next?”. Here are some situations I used with both Primary School children and 3rd-year undergraduates.



Figure 6. “What happens next?” Both scenes were purposely created with ambiguity. Scene (b) was composed to draw attention to the monochrome background and colored object, one of which is animate.

Both groups of children and students remarkably came up with the same answers. For (a) there was a group who said that the robin flew away; unfortunately, they did not notice the crutch so that the robin had lost the affordance to fly. A second group suggested that grog (green WeeBee) flew up and rescued the robin; unfortunately, they did not notice that grog doesn’t have wings, so did not have the affordance to fly. I’ll leave what the third group said to your imagination. There were two groups in (b), the first said that Miss Goble (teacher) sat on the rug and had a picnic and the second said that she took a ride on a magic carpet. Clearly there is scope to imbue objects with fictitious as well as real affordances.

A second investigation is shown in Fig.8 where a number of backgrounds were concocted, and some characters and objects to place in them were concocted. Participants were asked to one object with a background of their choice to create meaningful scenes. Here are some results (Fig.7).







Figure 8.



Towards a Theory of Affordances

In my research I found it useful to produce some theory of affordances, or more correctly how to classify these to allow me to create scenes with meshing. This theory is now quite old and needs enhancing, but here it is, since it to hand. It takes the form of two diagrams. The first is a *Systems Network* 'SysNet' which you read from left to right to get progressively more detailed information, Fig.9.

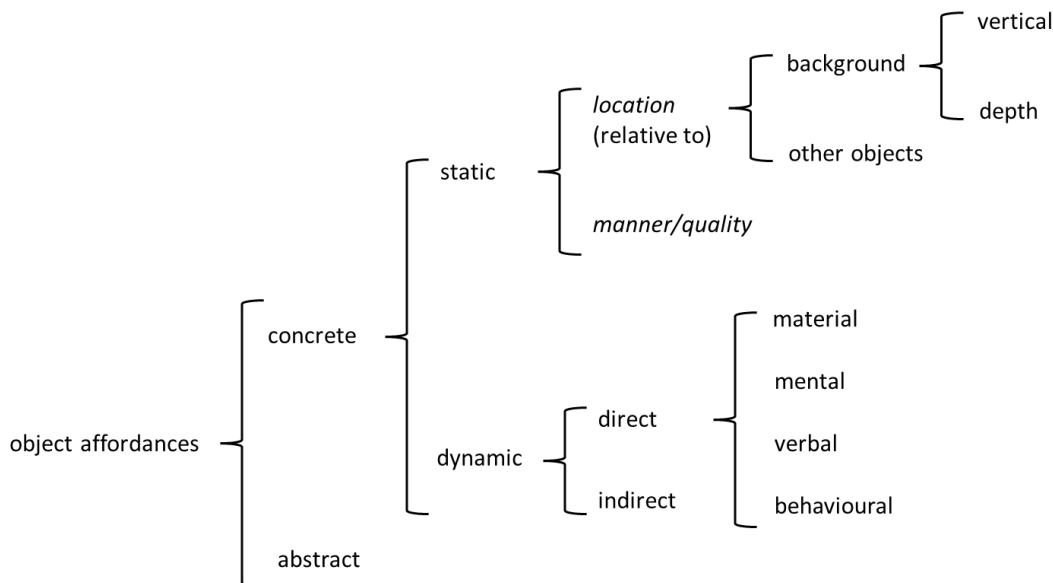


Figure 9. Example of a SysNet to help classify affordances.

Working from left to right, the first distinction is abstract and concrete, e.g., the abstract for Bob could be 'friendly'. A dragon character would be 'dangerous', and Miss Goble could be 'sweet'. This affordance trace is laid down in our memory by experience, with the animated Bob cartoon and with the real Miss Goble. Then we distinguish between dynamic (movement) affordances such as the ability of Bob to walk, and of Flup (a WeeBee with wings) to fly, and static affordances. Static affordances can be classed as either to do with location ('The boat is sailing in the lake') or manner/quality ('The rock is hard, sharp, medium sized and brown'). Again, all of these affordances are laid down as traces in our memory by experience with the real world. You can work out the rest of this branch of the SysNet.

Now to the 'dynamic – direct' part of the SysNet. These are informed by *linguistic processes*; material processes are concerned with action (running, hiding,...) mental processes are concerned with thinking and reflecting, verbal processes are to do with speech and behavioural processes are to do with extended actions, such as dancing, preparing a meal, a night at the opera.

In addition there is another dimension which we saw in the original Bob example, the difference between what Bob can do and what can be done to Bob; these are *operant* and *receptive* aspects. Consider a ball, it's operant attributes are what it can do, rest on a table, roll off the table, and its receptive attributes are what can be done to it, picked up, thrown but not anaesthetised! Compare this with a cat.

The second way of summarising attributes is in a table, remember Table 1. Here's Table 2. summarising the attributes of the WeeBee Pip:

Object
Visuals
Usual place
material
mental
relational
verbal
behavioural
Abstract
intransitive

Pip	
Pink, ref	
On the Ground	
Proc: Operative	Proc: Receptive
move to pick up ~fly	be moved be picked up
think, hear,	be thought about
Is on, becomes sad	
Talk, listen, think	
dance	
friendly, female	
is sitting.	



Table 2. Attributes identified for the character Pip

My experience with this sort of thinking has taught me that some of the most important words in designing scenes and animated stories is the class of words known as *prepositions* (where objects are). I think this may be of particular interest for this project. Here is the rather rich SysNet I ended up with (though I was asked to remove it from the journal article I got published).

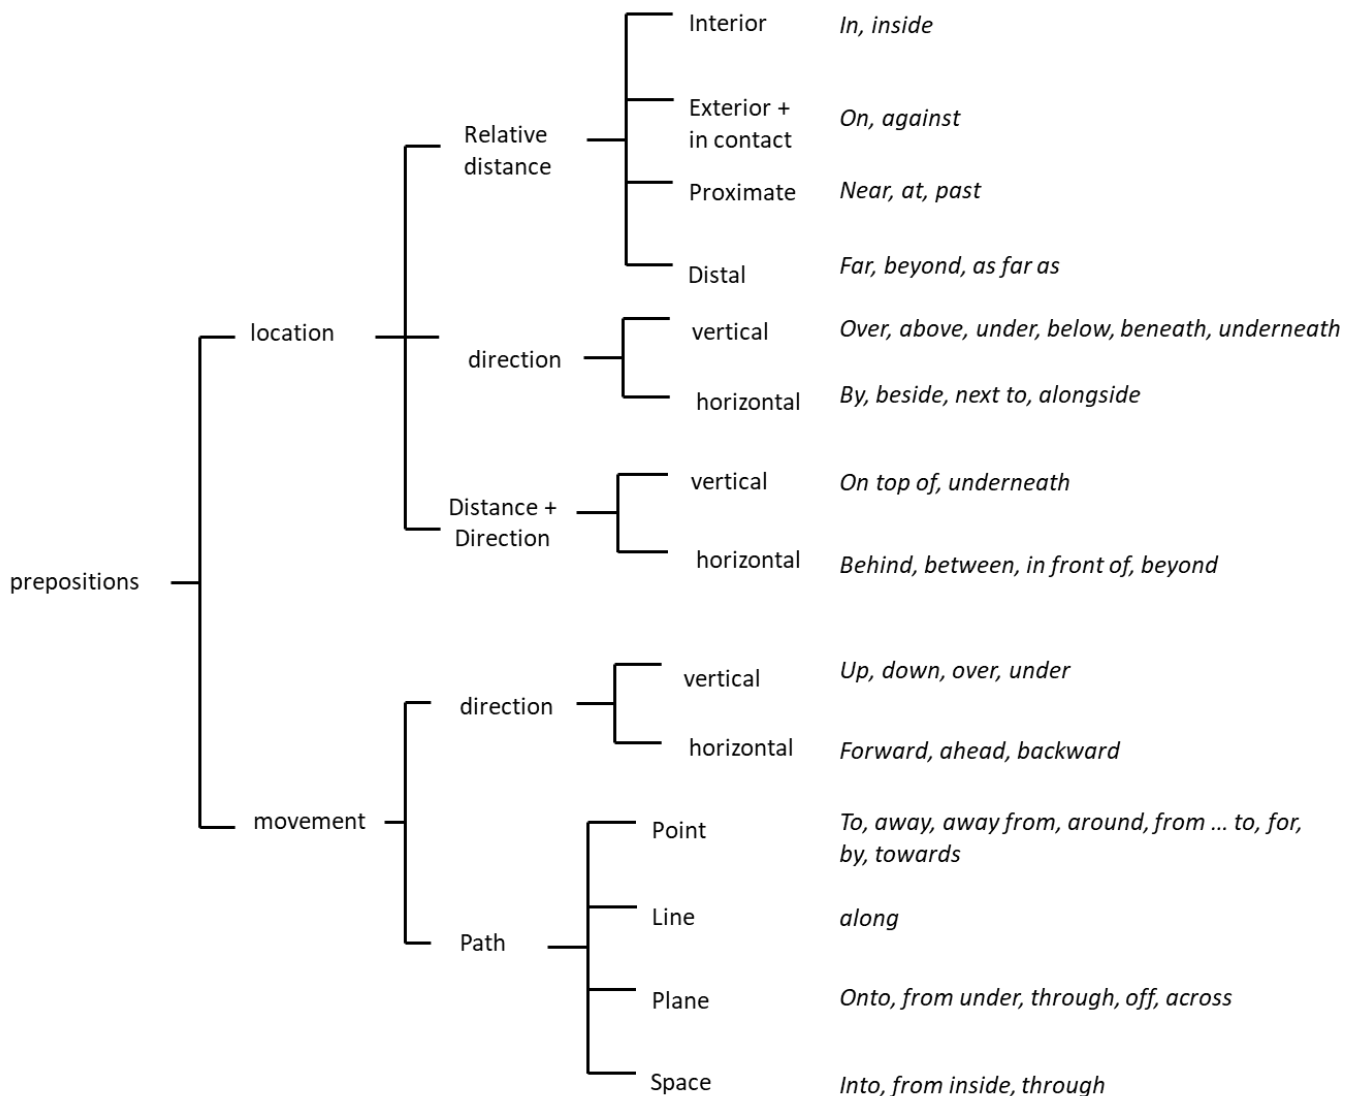


Figure 10. SysNet of prepositions with some examples.