

Can Agents Without Concepts Think? An Investigation Using a Knowledge Based System

Nicky Moss*, Adrian Hopgood** and Martin Weller*

*The Open University, Milton Keynes, MK7 6AA, UK

**Faculty of Computing Sciences & Engineering, De Montfort University,
The Gateway, Leicester, LE1 9BH, UK

n.g.moss@open.ac.uk

aah@dmu.ac.uk

Abstract

Grid-World is a working computer model which has been used to investigate the search capabilities of artificial agents that understand the world in terms of non-conceptual content. The results from this model show that the non-conceptual agent outperformed the stimulus response agent, and both were outperformed by the conceptual agent. This result provides quantitative evidence to support the theoretical argument that animals and pre-linguistic children may use non-conceptual content to understand the world. Modelling these ideas in an artificial environment provides an opportunity for a new approach to artificial intelligence.

1 Introduction

There does not seem to be a universally agreed definition of intelligence [1]. There are, however, some general characteristics that can be attributed to intelligent behaviour, and when witnessed in a particular situation a judgement can be made about whether a given agent is intelligent or not. One way of explaining thinking is to refer to concepts. Many philosophers [2] believe that only language users can have concepts, and along with this a view of the world that is objective. Others [3] argue that conceptual content has developed out of non-conceptual content, and that there must be some sort of objective understanding as a pre-cursor to developing into an agent with full conceptual capabilities.

It is the reflective nature of intentions which provides the mechanism to enable agents to think about their actions in ways that can lead to goal directed or purposeful behaviour. Intentions also form a key element in the belief, desire and intention (BDI) architecture and follow on from Bratman's [4] systematic framework for characterising mind and actions in terms of intentions. The relationship between beliefs, desires and intentions is easy for us to comprehend as it reflects the way we reason about the world in our own conscious minds. It is a

perfectly reasonable, and largely practical, way to construct an artificial agent. However, human reasoning is a conscious process that is intrinsically tied into our conceptual understanding of the world. For this reason most of our practical reasoning will use conceptual content, hence the success of the BDI architecture as an AI technique.

The approach adopted here is one that is based upon the multi-disciplinary principles of Cognitive Science. The idea that there is an intentional capacity that can be explained in terms of non-conceptual content is important because it frees us from linking an intelligent understanding of the world to agents that possess language. One difficulty with the general progression of non-conceptual content is that it has been largely based on discussion in the philosophical literature. Although there is plenty of supportive evidence from experiments using infants and animals, there have been few attempts to use computers to model non-conceptual content. Grid-World provides an environment that enables the characteristics that embody non-conceptual content to be modelled and investigated using artificial agents.

2 Non-conceptual Content

Our understanding of concepts is so dominant in our view of how things are in the world, that it is very hard for adult humans to think of anything in the world that is not defined conceptually. Holding concepts is clearly a conscious thinking activity, so it is the one that we are very aware of. Conceptual content is therefore about perceptual beliefs; it is how we believe the world to be. Non-conceptual content, on the other hand, is about perceptual experience, it is how the world is presented to us. An analogy can be drawn with the difference between analogue and digital. Information in the world is basically analogue in nature. For our minds to be able to grasp and manipulate concepts it needs to abstract away from the detail, and one way to do this is to encode it digitally with a certain level of quantisation.

Non-conceptual content can be defined [5] as a mental state that represents the world but which does not require the bearer of that mental state to possess the concepts required to specify the way in which they represent the world. According to this definition, it is possible for an agent to act as if it holds a concept, when in fact it does not. For example, an agent may search a problem space as if it has the concept of planning, but upon further investigation it may be revealed that the agent did not hold this concept at all.

Agents that are intentional will use intentions to work towards their goal (mean-end reasoning), maintain their intentions unless there is good evidence not to, constrain the range of options they may consider (constrain future deliberation) and allow them to change their beliefs (reason practically) as they discover new features of their environment. The ability to hold an intentional attitude has long been attributed to concept holders. Non-conceptual content provides an explanative mechanism that allows intentions to be attributed to non-language using agents.

Affordance has been established as the term used most frequently to reflect Gibson's ecological view of sensation, first put forward by him in 1950. Affordances offer a good mechanism for capturing information about the environment, and one that is sufficiently rich to provide the basis for action according to a non-conceptual view. If affordance can provide the opportunity, then non-conceptual content provides the mechanism to turn perception into action. The most immediate source of affordance comes from the agent itself, especially the role that the agent's own body plays in perception. This role of affordance in perception may be crucial for implementing artificial agents.

Bermudez draws upon some experimental work from psychology concerning object permanence. An agent is said to hold the idea of object permanence if it believes an object exists, even when that object is not being directly perceived. The basic experimental methodology is to place objects in full view of a very young child and then remove the object from view (usually by obscuring it behind a screen). The child is then tested to see if it appreciates the fact that the object could still be there. In other words to test the child's conceptual understanding of objects, and in particular its understanding of object permanence. The general conclusion has been that object permanence only comes with mastery of the concept of objects. Bermudez argues, by the following alternative explanation, that this need not be the case. It is possible for the child to be 'aware' of the object, for example, by showing surprise when it apparently moves through a solid screen, yet the same child will not search for the object behind the screen. This seems to provide some evidence that the child has representational capability, with respect to the object, even when that object is out of view, however, the same child obviously does not have the concept of an object.

3 Grid-World

Grid-World has been built using an expert system toolkit called Flex¹ (and associated prolog compiler) which uses its own English-like knowledge specification language (KSL). Flex supports frame-based reasoning with inheritance, rule-based programming and data-driven procedures. Grid-World makes extensive use of rule-based logic and uses frames to organise data. The features of Flex make it an attractive option for investigating agents that have basic sensory and memory capabilities.

Grid-World has drawn upon some of the ideas behind Brooks' subsumption architecture [6]. Brooks' general approach is to build robots with relatively simple rules and see what sort of behaviours emerge as a result. Grid-World also uses relatively simple rules, and behaviours do emerge. This bottom up approach is also important in the principal relationship between non-conceptual and conceptual content, as it can be argued that creatures that hold concepts may have evolved from

¹ Developed and maintained by Logic Programming Associates (LPA) Ltd.

others that had simpler capabilities. However, Grid-World falls well short of Brooks' approach where real robots are used in the real world.

Three different agents are implemented in Grid-World. The first agent – called stimulus response (SR) – searches using a combination of wall following and random searching and is modelled on the principles of stimulus response conditioning. The second agent – called non-conceptual content (CC) – is given the capability to recognise and remember features of the environment it has previously learned. It is then able to use these to aid any subsequent search. These capabilities model a form of non-conceptual content that uses affordances to recognise places. The third agent – called conceptual content (CC) – also has memory and place recognition capabilities, but it can also form a plan to join the affordances together. The ability to plan is a conceptual skill. The three agents can be viewed as steps from the simplest to the most sophisticated, and in each case the capabilities of each agent build upon those of its predecessor, to ensure that this reflects development rather than diversification.

The simple aim for each of the agents in Grid-World is to find its way from home to the goal. This choice of task, within an artificial world, has been chosen because it mirrors the actions of simple animals. This idea follows the work of Campbell [7] who suggests that basic navigation skills are ones that could be explained using non-conceptual content. Picking up on Gibson's [8] ecological approach to vision, Grid-World has affordances built into the environment, and these can be used by the two more sophisticated agents in helping them to reach their goal when searching. So the overall design principle of Grid-World is one that draws upon comparison with real world agents, and this principle is reflected in both the architecture and the rules that govern the behaviour of the three types of agent.

3.1 Grid-World Environment

The name 'Grid-World' was chosen because the space occupied by the world is defined in terms of x and y co-ordinates, where each co-ordinate defines a single square within the world. For example, the co-ordinate (4,5) is a single square 4 spaces to the right (x) and 5 spaces up (y), taking the co-ordinates (1,1) as a point of reference that represents the bottom left hand corner of Grid-World. So the whole world can be seen as a grid of intersecting lines, with the spaces between the lines defining a large number of squares. Each co-ordinate contains either a blank square, or a tile representing some feature of the world as will be described below. A visual representation of Grid-World, as appears on the computer screen when the program is running, is shown in figure 1.

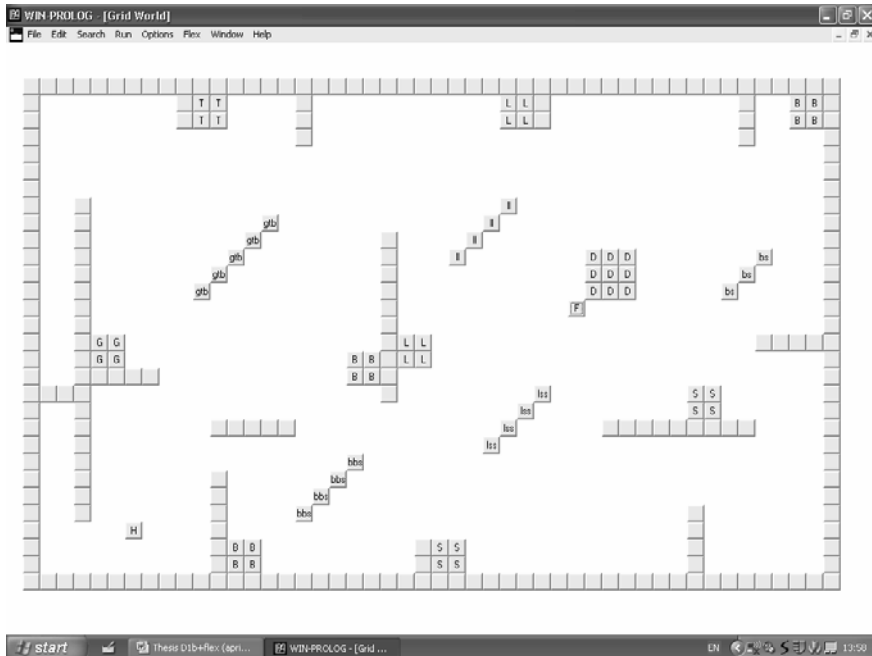


Figure 1. Screen view of Grid-World

The entire space is bounded by a continuous wall that occupies all the outer grid squares. Each square that makes up a wall contains a blank tile with no labels. In the example shown this is a boundary measuring 48 by 30, giving a large internal space of 1440 squares. Horizontal and vertical walls are added to introduce features into what would otherwise be an open space. These are also shown as blank tiles. A goal is defined as a block of one or more squares, and each one is marked with the letter D.

Other features of the landscape are also shown as blocks of tiles, each labelled with a letter. These represent various landmarks. For example, the block labelled 'L' is a lake, and the block labelled 'T' is a group of trees. The main purpose of the landmarks is to provide some visual cues, things in the environment that the agent may recognise, and for this reason the agent is able to pass through the landmarks. The blocks that have two or three lower case letters represent areas of Grid-World that lay between two or more landmarks. These are called affordances as they provide the agent with information directly about the environment. One way to think of them is as points in the landscape where one or more distant cues intersect. The tile marked with an 'H' is the home location and is where the agent starts its search. The position of the agent at any given point during a search is shown by a tile with the letter 'F'. This was chosen to represent forager, and this tile moves as the agent searches the landscape.

Grid-World has been designed to provide the same sort of complexity, for an artificial agent, as experienced by a small mammal, for example. The walls are primarily to provide some features that can be used by the stimulus response agent, recognising that wall following is a basic search strategy used by many animals. The walls also provide some variety to the space, making searching a little more complicated for all agents. The cues, and associated affordances, are intended to be used by the non-conceptual and conceptual agents and mirror some of the features used in experiments where mice or rats search mazes.

3.2 Stimulus Response Agent

The basic search characteristics used are a combination of wall following, object avoidance and searching open space. An agent moves one square at a time. If it finds that a square is occupied by a wall tile it will not enter that square. The agent will recognise a wall whenever it finds a continuous string of wall tiles. Whenever it finds a wall it will follow it until it reaches the end. The open space search involves going in one direction for a number of squares, before turning left or right. The number of squares that an agent moves in any one direction is chosen randomly from a preset range, and the choice of left or right alternates. This basic strategy ensures that all squares can be searched through a combination of wall following and the random searching of open space. With this basic strategy the goal is always found. The more sophisticated skills needed for the agent to search intelligently are built upon these basic characteristics. All agent types will fall back to a random search when there is no better choice.

The agents location, and the location of all the Grid-World features, are stored in a range of databases. Information is read to and from these databases using actions of the format:

```
read_changes (position(X,Y,T))
```

In this example position is the name of the database, x and y are the location coordinates and T is the data being read.

A combination of rules and actions are then used to move the agent according to the information retrieved. For example, an action that would start an agent searching open space has the format:

```
Action random_search_left;
  do n := 0
  and random_number(12)
  and while in_open_space
  and n < random_range
  do move_left
  and n := (n+1)
  end while .
```

The agent continues searching until it finds the goal. The program then records a range of quantitative data including the number of steps taken to move from home to the goal and a list of the affordances visited.

3.2 Non-conceptual Agent

The non-conceptual agent learns during the initial search phase which can follow the same path as the stimulus response agent. During this phase, whenever the agent finds a point of affordance, it makes an entry into a memory frame. The principle behind the memory is for the agent to record the direction of the next point of significance, either another affordance or the goal. For this purpose it makes a note of its current location and then records the net direction when it reaches the next point of significance. This memory also includes the name of the affordance at which it was initiated.

An instance of a completed frame is created whenever an affordance is found during an initial search. Each instance is given a number and records a start point (starting), an end point (ending) and the direction from start to end. The direction is simply recorded as a compass bearing that divides the area around any square into four quadrants.

In the non-conceptual search mode the agent has access to the additional information held in the frames created during the initial search. It uses this data whenever it finds an affordance. The principle behind the agent that searches using non-conceptual content is that it gains a benefit from understanding the significance of the affordances found in the initial search. This significance comes from recognising the place and knowing something about the last time it visited this place. What it knows is that setting off in a particular direction moves it towards another place it would recognise. It does not know that this may eventually lead to the goal it just recognises that it did set off in a particular direction last time, so it chooses to go this way again. The non-conceptual agent gains this benefit from reading the frames made during the initial search.

All other aspects of the search remain the same and the agent still follows walls, for example. To ensure that the agent is given no other advantage, apart from setting off in a particular direction, when it sets off in the remembered direction it still uses the same range of random numbers to move as it does when it is in open space. One way to think of the advantage that this agent has is just being 'nudged' in a potentially successful direction.

The consequence of the data in the memory frame is to point the agent in the general direction of the next affordance and set it off in that direction. For affordances that are near the goal, it is possible that the agent will be pushed towards the goal. There is no guarantee that any instance will point the agent towards another affordance closer to the goal. If, during the initial search, the agent happened to find an affordance that was further away from the goal, then this is the direction it records, and this is the one the non-conceptual agent will use.

3.3 Conceptual Agent

The last search phase, as used by the conceptual agent, is modelled on an agent that holds navigational concepts. Most significantly it is able to take a detached view of what it has done and plan accordingly. The difference between this search strategy and the one used by the non-conceptual agent is best understood by looking at two aspects of that search that may not be optimised. During the initial search phase the creation of memories is dependent upon the agent finding an affordance and moving from that affordance to one that takes it closer to the goal. It is possible (although unlikely) for the agent to find the goal during an initial search without finding any affordances. Any subsequent non-conceptual search would be identical to an initial search as there is no additional information to help the agent. It is more likely that the agent will not create a memory for each affordance type as it will not visit them all. It is also possible that an agent will move from one affordance to another that is further away from the goal. Neither of these will optimise the agent's chances of finding the goal.

Because the conceptual agent can take a detached view of the search space, and it is capable of planning, it can be given the optimum memory for each of the affordance types before it searches. This plan is created by having a frame for each affordance that always points the agent in the direction of the goal. The agent then uses the same search strategy as the non-conceptual agent, but one that draws on this optimal memory and not what is learned by any other search.

4 Experimental Results

There are three agents. Each one has the same basic sensory and search capabilities. The three differ in the way they represent sensations that they detect in the environment. The first agent (SR) responds in a strictly law like way to any stimulus when choosing the next step. The second agent (NCC) uses sensory information that is immediately available to it, combined with memory to assist it in making a choice about where to go next. The third agent (CC) uses sensation together with a plan to choose the next move.

4.1 Experimental Design

The difference between each agent is the way in which they represent sensory data in their environment. Therefore, there is one independent variable; the type of representation. In this design there are three levels for the independent variable that correspond directly to the implementation of the three agents. The dependent variable is the data that records the number of steps it takes the agent to move from home to goal.

The experimental hypothesis predicts the relationship between the independent and dependent variable. In this case this is the prediction that the search performance

(dependent variable) will improve as the agent's search capabilities become more sophisticated (independent variable).

Statistical tests were applied to the results to determine whether observed differences in the mean were due to the experimental conditions or chance. A result was deemed significant if there was a probability of less than 5% that the result occurred by chance.

4.2 Results

The statistical data for 40 searches carried out by the stimulus response (SR), non-conceptual (NCC) and conceptual (CC) agents are shown in figure 2. In this experiment the SR agent followed 40 different paths, and each one of these paths were used as the initial search (training) for the NCC agent. The CC agent used the same pre-set memories for each search.

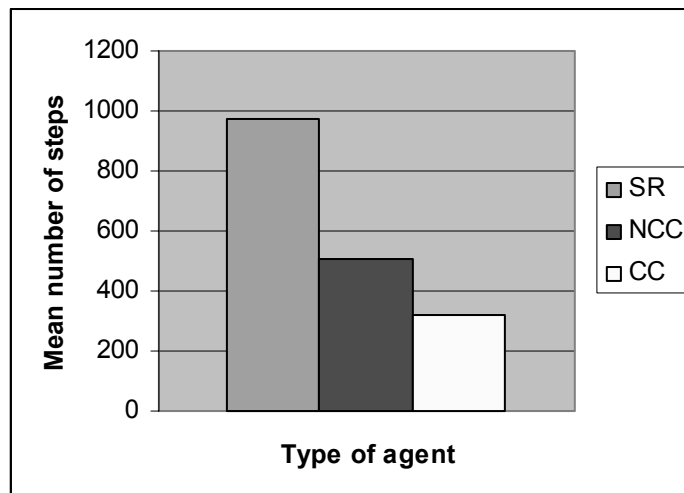


Figure 2 – 40 searches, with 3 types of agent

The mean values for the stimulus response, non-conceptual and conceptual agents are 1174, 419 and 290 respectively. These results show a significant difference in the search performance of the three agents.

By looking at the data for individual searches, aided by information on the distribution of step counts, it can be seen that all three agents manage to find the goal in fewer than 100 steps in at least one search. This represents a very good search performance, as the minimum path length is 40 steps. At the other end of the spectrum it can be seen that the stimulus response agent takes the longest single search time of more than 5000 steps. The fact that the stimulus response agent reaches both extremes is not surprising as any basically random search method is likely, over enough runs, to find the extremes. The overall distribution of searches

shows more clearly that the non-conceptual and conceptual agent are able to find the goal in fewer steps on many more occasions. This is particularly true for the conceptual agent which only once took more than a 1000 steps out of 80 searches.

While the results show that there is a clear difference between the search times for each of the three agents, the range of search times varies considerably for the stimulus response and non-conceptual agents. Other data also shows that the number of frames generated by the stimulus response agent also varies across a wide range. Figure 3 shows the total number of times that each affordance was used, for each of the three agents, during another set of 20 searches.

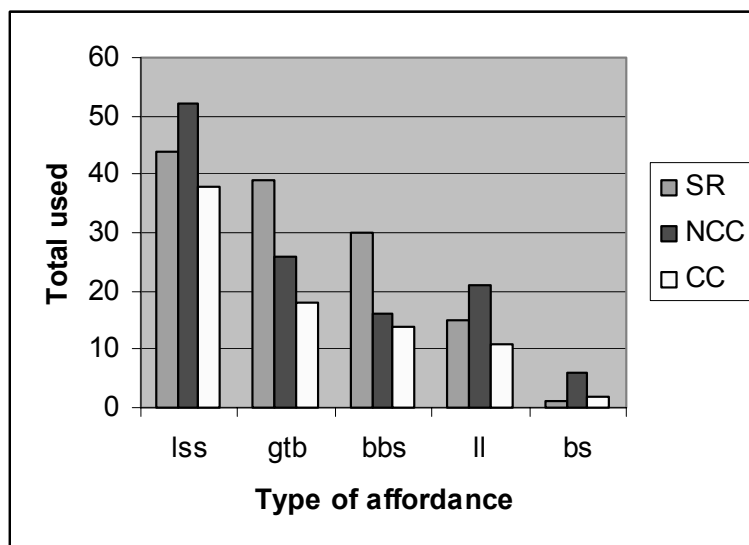


Figure 3 – How often an affordance is used

As can be seen in figure 1, the affordances are located in the ‘open space’ areas of Grid-World. These locations were chosen to make sure that the agent would have to find the affordances, and not be led directly to them by following the wall, or stepping directly into one when it left the end of a wall. The largest affordance (gtb) occupies five tiles, and the smallest (bs) three. These sizes were chosen because these were respectively placed in larger and smaller areas of open space.

The lss affordance is visited most often by all three types of agent, and bs is visited least often. The other three affordances are visited often, and normally occur in each search. Looking at which of the affordances is visited just before the goal is reached (the last affordance) shows that lss is very dominant in this role, occupying this position in just over 73% of the searches carried out in experiment 4. By comparison, bs only occupied the last position in just one search.

From these results, it would appear that lss is playing a dominant role in the search strategy used by all three agents. This is partly explained by the proximity of lss to

the goal, although it is not immediately obvious, from the results, or looking at figure 1, why lss should be more dominant than ll, for example. After all, both are close to the goal and on the same side as the agent's home. The fact that bs is situated on the far side of the goal, which means that the agent has to pass the goal to visit it, could be part of the reason that it is often not included in a search very often.

Other arrangements of Grid-World were used to explore the impact of different environments on the search performance of the agents. As a general rule, the non-conceptual agent outperformed the stimulus response agent, and both were outperformed by the conceptual agent. However, the location of affordances, both in absolute numbers and in their relationship to each other and wall, was significant in the effect upon both the non-conceptual and conceptual agents. Another set of experiments was also able to show that if the non-conceptual agent could string together affordances in the correct order, then it could match the performance of the conceptual agent. This result re-enforces the idea that the difference between the two is rooted in the latter's ability to plan.

5 Conclusion

The overall aim of developing Grid-World has been to provide a flexible experimental environment that can be used to explore non-conceptual content using qualitative data. Using Artificial Intelligence techniques in this way is consistent with a cognitive science approach where computer models are used to investigate theoretical ideas. Grid-World has provided a robust artificial environment in which it has been possible to implement different agents and to test their search capabilities. The implementation has also enabled the test environment to be altered to explore agent behaviours as they emerged from earlier experiments.

The results from all of the experiments using the stimulus response agent have shown that a combination of wall following and searching in open space was sufficient to always lead the agent to the goal. However, with search values ranging between 40 and 5,000 steps, there was considerable variation and this cannot really be called a successful search strategy.

The non-conceptual and conceptual agents both use a memory of affordances to give significance to proximal features of their environment that they sense. The location and size of affordances within the search space was shown to have a measured impact upon the search performance of both agents. These characteristics of affordances provide further evidence to support the idea that there is a real proximity limit to an agent's capacity to grasp the significance of an object. If it is too far, it is just beyond the agent's reach. So the distance between any two points of affordance must fit with the agent's capabilities, and when searching there must be a sufficient number of well spaced affordances to provide a path. This conclusion is supported by the actions of the conceptual agent, which even with a path linking affordances, still did less well when affordance were moved near the goal.

Only a very basic memory capability has been given to the non-conceptual agent. In particular the memory instance for each affordance is only significant to the agent at that unique location and it cannot access memories from one location at any other. Further, although the agent can visit an affordance more than once during an initial search, it always recalls the result of the last visit. So it cannot learn the best direction from many visits. All of these parameters were carefully chosen so as not to give the agent any concepts, yet the non-conceptual agent has still consistently found the goal in fewer steps than the stimulus response agent. In some way it has operated as if it had a concept of searching, not one as sophisticated as the plan held by the conceptual agent, but nevertheless one that might convince an observer that it held a concept of planning.

This is an exciting conclusion as there is little other experimental evidence to support the case for non-conceptual content. However, a more critical reflection tells us that much of the performance improvement is closely related to the positions of the affordances within a particular environment. This is not to deny that there is an agent in Grid-World that can be said to have non-conceptual content, just that it is necessary to appreciate the close relationship between such an agent and the affordances in its environment.

If the behaviour of animals can be explained in terms of non-conceptual content, then there are reasons to believe that conceptual content has evolved from this simpler representational capability. If this is the case, then the evolution of conceptual content from non-conceptual content provides evidence to support a bottom up approach to artificial intelligence.

References

1. Weiss, G. (1999). *Multiagent Systems: A Modern Approach to Distributed AI*. Cambridge MA: MIT Press.
2. Gunter, Y.H. (2003). *Essays on Nonconceptual Content*. Cambridge, MA: MIT.
3. Byrne, A. (2003). Consciousness and Non-conceptual Content. *Philosophical Studies* 113: 261–274.
4. Bratman, M.E. (1987). *Intention, Plans and Practical Reason*. Cambridge, MA: Harvard University Press.
5. Bermudez, J.L. (2003). *Thinking Without Words*. Oxford: Oxford University Press.
6. Brooks, R.A (1999). *Cambrian Intelligence – The Early History of the New AI*. Cambridge, MA: MIT Press.
7. Campbell, J (1994). *Past, Space & Self*. Cambridge, MA: MIT Press.
- 8, Gibson, J.J. (1986). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.