A Cognitive Model for Learning and Reasoning over Arbitrary Concepts

Kieran Greer, Senior Member, IEEE.
DistributedComputingSystems.co.uk, Belfast, UK.
e-mail: kieran.greer@ntlworld.com.

***Abstract* - This paper describes a new cognitive model that could be used to build a system that can intelligently realise concepts for itself and then reason over them. The main model has been published previously, but this paper proposes a new finer level of processing that would allow the system to learn arbitrarily complex concepts for itself. These can then be clustered into chains that represent higher level concepts and reasoned over. These chains can also trigger each other to generate a certain level of 'thinking'. This model would be suitable for a neural-like system, but also a large distributed network.**

Keywords – cognitive, distributed network, learn, concept, reason.

I. INTRODUCTION

This paper will try to tie together previously published work, suggesting how a cognitive model might be built that can actually 'think' [1][2]. This paper brings different components of the earlier research together, but then also tries to add a new finer level of processing to the model that will allow it to learn arbitrarily complex concepts for itself. This paper only proposes the model, with no results to justify it, but it is based on well known mechanisms and the underlying theories have been proven.

Section II describes the model in relation to a neural-like system with symbolic representations. Section III describes how solving the image processing problem will allow the model to learn arbitrarily complex concepts for itself. Section IV describes the new model with the new finer level of processing, while section V gives an example application of this model. Section VI discusses what the immediate future work might be and section VII gives some conclusions on the work that has been described.

II. SYMBOLIC NEURAL NETWORK

It has been suggested in [1] and [2] that it could be possible to build a neural-like structure based on symbolic elements or nodes. This means that it would be possible to create a neural network where all of the nodes have some sort of symbolic representation. Symbols are important because this allows the system to understand what each node means or represents. Newell and Simon [4] are noted for supporting a symbolic approach. With their 'Physical Symbol System Hypothesis' they state the need for a symbolic representation of the environment, so that a computer can understand it. However, a universal machine that can create, understand and use these symbols remains a problem to be solved. In order to do this there needs to be another level of intelligence above the symbolic level that can intelligently use the symbols. So while the symbols can provide suitable descriptions, extra intelligence is required to successfully manipulate them. But for this intelligence to be useful in a generic sense, it needs to be flexible and dynamic enough to deal with many different scenarios. The cognitive model suggested in [1] and [2] tries to combine the symbolic approach with a neural-like structure that might provide the required level of flexibility, to allow the symbols to be used in such a way that the machine itself can generate its own understanding. This model actually arose out of research that tried to optimise a distributed network with respect to query processes and also allow reasoning capabilities as part of the query process. The cognitive model that was developed is shown in Figure 1.

This model contains three different levels of intelligence. The first or lowest level allows for basic information retrieval that is optimised through dynamic links. The network can be queried using a query language and information that is directly identified can be retrieved. The dynamic links will then provide for query optimisation by reducing search time. The linking mechanism works by linking nodes that are associated with each other through the use of the system. It is based on the stigmergic process of linking through experience, but it could also be called Hebbian.

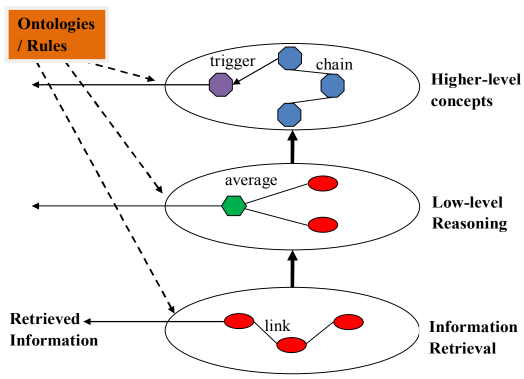

Figure 1. Cognitive Model with three levels.

This level of the model has been extensively tested and results show that it could reduce the number of nodes searched by 80 – 90% with a related depreciation in the quality of answer of only 5 – 10%. It has also been explained that the linking process is completely generic, relying only on the feedback of associated nodes. It is thus not tied to any particular type of information and so could link any nodes that have a suitable symbolic representation. Detailed summaries of the testing of this level can be found in [1] or related papers, which are all available from the 'licas' web page [3].

The second or middle level allows for low-level reasoning, where questions such as 'what is the best value for one source based on other source values' or 'is a value or action possible based on other values' can be answered. These are averaging queries (best) or queries to test if a condition simply exists. The dynamic links typically represent the knowledge of the users of the system, in the form of the associations that they make between the individual nodes or concepts that they query. If the users' queries are reliable, then the links that are formed can be used for more knowledge-intensive querying. For example, users may typically query about 'city attractions' and 'transport schedules', but they would not query about 'city attractions' and 'the latest mobile phones'. Through simple mathematical operations, such as the averaging or aggregation of links to the same data types, low-level reasoning over the knowledge can be performed. For example, if there are several links to restaurants that have been visited by users of the system, then averaging over these, in effect selecting the most popular restaurant, could give a general assessment of what the best value would be. There has also been some testing of this level with positive results. This would be expected however, as the general architecture would naturally prune worse nodes and link better ones, biasing any averaging process.

The third or upper level represents higher-level reasoning, where the network autonomically generates more complex concepts, allowing for more complex reasoning to be performed. To do this, the source concepts that typically occur together need to be grouped into higher-level chains of concepts. These higher-level groups can then be used for more complex reasoning. It may even be possible to code rules into the network, where one concept or group would automatically trigger another one. If this can be achieved, then the network can begin to realise concepts for itself and thus to reason or think for itself. This is the level that is of interest in this paper and may be particularly applicable to a cognitive model using a neural network-like structure.

The final architecture also incorporates a knowledge-base, with ontology and rule-based information that can be used at all levels. The dynamic links that form the main reasoning mechanism can then be used as follows:

1. *Higher-Level Reasoning* - Dynamic links associated together, forming chains of higher-level concepts. One chain or concept can trigger another chain, which if then realised could trigger another one and so on. The network then begins to realise these higher level concepts for itself, which is a kind of thinking.
2. *Low-Level Reasoning* - Dynamic links used directly for aggregation/averaging queries.
3. *Direct Information Retrieval* - Dynamic links used to optimise the network and query process.

III. IMAGE PROCESSING PROBLEM

As a bit of fun, the book [1] concluded with a proposal for how a computer might also use image processing to generate internal understandings of its environment, but this idea can also be extended further into a more complete model. Human beings are able to build up internal pictures of the world around them that is not all represented by words and numbers. Blind people are also able to build up a picture of their environment without the help of sight. So we also have an internal vision that we use when thinking. Most AI systems place certain restrictions on the computer, with regard to the machine's freedom to process the information that it is fed in any arbitrary way. For example, logic-based languages could be used to describe everything, but they are still quite restrictive in the way that they represent things.

The new idea is to give the computer much more freedom in being able to describe its environment. For example, an object might be described to the computer, which would then be asked it to 'draw' what it understood about the input. It would then need to have an internal conceptualisation of the environment and be able to communicate this to the program user. For example, if a table is described to the computer, it then might try to draw the table top, but would need to know that this is represented by four lines in a rectangular shape. It would have to be able to draw each line individually and know that they must join up. It would not be allowed to represent the table top as a single concept of a rectangle that could be retrieved from some library. Because the feedback mechanism would be less formal, it might also be more interesting and yield insights into what the computer actually understands. The computer program could also try to process the images internally using some image processing algorithm. This might allow it to extract more subtle similarities or differences in the different input that it is being fed, rather than the explicit descriptions of a formal language. Then it might be able to learn some 'tacit' knowledge for itself.

The advantage of this is that the computer could try to learn or understand 'bits' of an image or concept that could not easily be described explicitly, such as one table with rounded corners and another with angular ones. This is because the image is no longer represented as a whole, but by many different related lines or curves, etc. It then might be able to feed these bits of information back to the researcher, if it is given a flexible medium in which to do so, such as a free drawing.

IV. NEW MODEL

One problem that was noted with the current model is that each concept needs to be symbolically tagged and this could be difficult for a computer to do automatically. For example, if a computer is processing an image autonomously, how would it know what tags to give the different elements that it sees? As described in [5] 'The human brain is composed of approximately ten billion cells, called neurons. These cells interact by means of electrico - chemical signals through their synapses. Even though there may not be very many different types of neurons, they differ in the structure of their connections'.

So this suggests that the human brain does not tag each neuron with a different symbolic meaning, but rather it is the pattern of neurons that all fire together that defines what the brain is currently thinking. This collection of neurons could then be interpreted by a more centralised component that can see what the collection actually represents. The collection could be compared to the chains of concepts in the model of Figure 1, the difference being that the chained nodes in that model are all tagged with a symbol.

It might however also be possible to look at a collection of associated neurons that fire together as an image of some sort. For example, an image displayed on a computer is made up of a large number of pixels that are associated with each other in some way. Each pixel is no different than another one and it is only our ability to recognise the patterns in the pixels that allows us to understand what the image is. A pixel collection would be the most flexible way of describing an image. So the collection of firing neurons could possibly also be represented in a similar way. If the system consistently received images that were similar in some way, then it could use these to represent something. This could add a new finer layer of representation to the model, where each symbolic concept is actually a collection of even finer grained elements that have been shown to be related to each other. Each new concept would then be a different pattern of elements. These individual concept patterns would then link together to form the chains of higher level concepts that can then trigger other chains, and so on. The symbolic meaning is possibly then not so important, because if it is possible to recognise the individual pixel-level patterns and have them associated with each other, this is what the system will understand and use to determine its next action. It would only require the symbolic tags when trying to describe these patterns and associations to somebody or something else. Note that the number of neurons in the human brain indicates just how fine grained its processing capabilities are and is still out of the range of current computer processing capabilities.

V. EXAMPLE APPLICATION SCENARIO

The most obvious example of how this sort of system might work is in fact through the processing of images. The system would receive input, from video or picture images, represented by pixels. These pixels can then be processed and filtered to realise certain concepts that are then linked together into chains that represent higher level concepts. The system does not have a symbolic name for these pixel collections, but if it then sees the same pattern again, it can recognise it and retrieve the higher level chains that it belongs to. The system can then begin to make higher level associations autonomously, or by itself. Because this is at a very small level of granularity, it is a very flexible way to represent the input, almost arbitrary. This would thus mean that there are no restrictions in what the system could learn.

This kind of grid-like structure could also apply to a sensorised environment. An obvious application is to monitor the movements of large collections of people, as in a military game-playing scenario. One collection of people could be recognised as one concept. If it then teamed up with another collection the new combined group might do something. The network would learn the association between these two concepts and then trigger the new action that occurred when they joined. Each sensor reading would be mapped onto a grid element, where each element would retrieve the sensor id and value from the sensor; while time and location would be stored automatically. At one moment in time we have one particular pattern, which could represent a single concept or several concepts. Pattern processing and filtering would identify individual concepts and associate them with each other to form concept chains. The next time period gives a different pattern that requires the same kind of processing and then the differences in the two patterns or images can be determined. The consistency or difference in the signals received determines how the pattern is stored or used.

Figure 2 gives an example of how this might work. The left-hand side shows the grid structure with activated sensors on each square. The right-hand side shows the concepts that are generated in the network. Colour-coding relates the concepts to the sensors. The top level pair show the initial situation, where the red and blue concepts have already been identified. The middle pair shows a filtering process that realises the two new concepts (green and brown) and the movement linking the first two concepts. The bottom pair then suggests that the first association has triggered the second one.

This research originally arose out of looking at optimising a network with regard to query processes. It was quickly determined that connecting two nodes that were often associated with each other directly was not accurate enough. What was needed was a descriptive path related to the link, to make the link description specific enough to be accurate. For example, the query process testing optimised an SQL-like query including the source and value types as part of the descriptive path. However it was also required to include the comparison operator to make the linking process accurate enough. A query such as 'Select A.Value1 From A, B Where A.Value2 GT B.Value3' would create a link between the A and B source nodes described by a path with the elements:

Source: B_source_node,
Key: Value3 - A – Value2 - GT - A_source_node

With the proposed model of a grid-like structure representing time and space signals, this path of concepts is not obviously available. In the case of a sensorised environment, then metadata describing each sensor could possibly be used to build a path of related concepts.

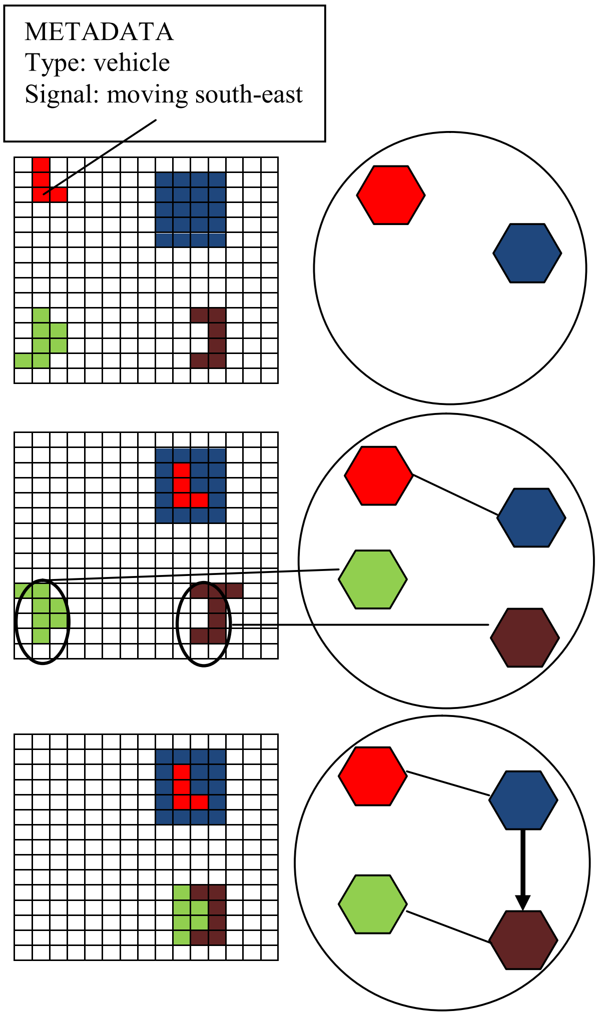

Figure 2. The formation of concepts and dynamic associations through pixel-like processing.

The problem would be in deciding what descriptive tags should be used to make up the key path. Alternatively, if the granularity is fine enough, then possibly the number of different signals that make up a concept would be specific enough to provide the desired degree of accuracy.

VI. FUTURE WORK

This model as a whole is still only a research proposal and there is a lot of work required to determine if it could work. While the first two levels of Figure 1 have been tested and returned positive results, future research would have to prove that the highest level is also feasible for the whole model to be valid. The first stage would be to try and solve the right-hand side of Figure 2. This would initially be tested with previously defined whole concepts. The system would be presented with arbitrary groups of these concepts, retrieved from consistent underlying patterns and try to learn the correct associations between them and how they trigger each other. If it can be shown that the system can in fact associate the correct concepts together autonomously and also learn the rules that would trigger one chain from another, then these principles could also be used to try and learn the finer grained arbitrary concepts.

VII. CONCLUSIONS

This paper has described recent research that has suggested a new cognitive model based on dynamically linking related nodes in a neural-like system. These nodes represent concepts and the links represent chains of higher level concepts that can be reasoned over. The new model then applies the same approach at a finer level, to try and create each individual concept in the same way. This would allow for the learning of any arbitrary concept, giving maximum flexibility and generality to the system. The main problem would be to be able to autonomously define the correct key path tags. The advantage with this method however is that at least the system would have enough information through the path tags to generate accurate enough links. This might simply not be possible linking two concepts directly.

REFERENCES

[1] Greer, K. (2008). "Thinking Networks – the Large and Small of it: Autonomic and Reasoning Processes for Information Networks", published with *LuLu.com,* ISBN: 1440433275, EAN-13: 9781440433276.
[2] Greer, K., Baumgarten, M., Mulvenna, M., Curran, K and Nugent, C. (2008). "Autonomic and Cognitive Possibilities for Information or Neural-Like Systems using Dynamic Links", WSEAS Transactions on Systems, Issue 9, Vol. 7, pp. 777 - 792. ISSN: 1109-2777.
[3] Licas, http://licas.sourceforge.net (last accessed 30/8/09).
[4] Newell, A. and Simon, H.A. (1976). "Computer science as empirical inquiry: symbols and search", Communications of the ACM, Vol. 19, No. 3, pp. 113 – 126.
[5] Weisbuch, G. (1999). "The Complex Adaptive Systems Approach to Biology", Evolution and Cognition, Vol. 5, No. 1, pp. 1 - 11.