Modelling Human Intelligence: A Learning Mechanism

Enrique Carlos Segura[1] and Robin Whitty[2]

[1] Departamento de Computacion,
Facultad de Ciencias Exactas y Naturales,
Universidad de Buenos Aires,
Ciudad Universitaria, Pabellon I, (1428) Buenos Aires, Argentina
[2] School of Business, Computing and Information Management,
London South Bank University,
London SE1A 0AA, UK

Abstract. We propose a novel, high-level model of human learning and cognition, based on association forming. The model configures any input data stream featuring a high incidence of repetition into an association network whose node clusters represent data 'concepts'. It relies on the hypothesis that, irrespective of the high parallelism of the neural structures involved in cognitive processes taking place in the brain cortex, the channel through which the information is conveyed from the real world environment to its final location (in whatever form of neural structure) can transmit only one data item per time unit. Several experiments are performed on the ability of the resulting system to reconstruct a given underlying 'world graph' of concepts and to form and eventually maintain a stable, long term core of memory that we call 'semantic' memory. The existence of discontinuous, first order phase transitions in the dynamics of the system is supported with experiments. Results on clustering and association are shown as well.

Keywords: Association network, memory, learning, graph, stability.

1 Introduction

The problem of building a model of human intelligence requires the distinction between two aspects: 1) storing information (learning) and 2) thinking about this information (cognition). However, it is not plausible a separation between these two functions in the model itself, neither physically (implementation) nor temporally (a learning phase followed by a cognition phase). Associative or distributed memory models [5],[6] remove this separation but at a low level of individual 'concepts'. We propose a model at a higher level: a collection of networks continuously stores incoming data from the outside world.

Suppose we provide a tourist with a list of the London bus routes, but without information on the topology of the town (e.g. a map). He or she can only pick different buses and concatenate pieces of routes so as to make up a continuous tour. Although the process of constructing a mental representation of

A. Gelbukh, A. de Albornoz, and H. Terashima (Eds.): MICAI 2005, LNAI 3789, pp. 1–10, 2005.

that topology is a complex, high-dimensional one, we can assume that he will ultimately be able to do it from the temporal sequence of inputs (bus stops, say).

In this paper we present a system, the Association Network (AN) that is able to model this and other situations involving some learning process from a certain real environment, and confirms the ability to infer complex topological relations from a single input stream, as would be expected from an intelligent agent.

The example above illustrates the philosophical basis for our model:

> learning = the reconstruction of some network on the basis of random walks on that network.

The network is a hypothesised worldview or 'world graph'; its reconstruction, incorporating approximations, abstractions and lacunae, may be thought of as long-term memory. How do we learn our way around a new neighbourhood? By making random walks in it. How do we learn natural language? By random walks on some hypothesised language network. How do we learn a game of chess or football? By random walks in the network of possible sequences of play. Of course, these walks are not Markovian (although our London bus routes were traversed in this way) but rather are generated by some stochastic process. The human brain is thereby presented with a sequence of random inputs and must configure these to reconstruct a worldview. This is done by making associations via repetitions in the random walk: it is striking how much structure can be imparted to AN's, merely by local identifications based on such repetitions.

We structure the paper as follows: in section 2 we present the basic mechanism for storing an incoming data stream into a network. In section 3 we present experiments upon randomly generated worldview graphs and upon a worldview graph of London bus routes. These experiments are used to confirm that interesting and stable structure may be stored in the network model. Finally, in section 4 we discuss our model in the context of AI and interactive computing.

2 Association Networks: The Learning Algorithm

Our network will be built according to four principles:

1. input data is stored in the nodes of the network;
2. at any point in time, in a non-empty network, there is a **current node** with which any new input will be associated;
3. there is a notion of distance in the net; and
4. there is a **threshold distance**, such that two nodes with the same data and being at most this distance apart, are to be identified with each other.

We may take the distance between two nodes to be the smallest number of edges in some directed path from one to the other. In our example we will set the threshold distance to be two edges. Assume the network is initially empty and suppose that the following input stream is presented to the network:

$$a\ b\ c\ a\ d\ c\ e\ a\ c\ e\ c\ e\ a\ f\ c\ a\ c\ b\ a\ c\ d\ a$$

The network starts off as a single node, which we draw with a square to indicate that it is the current node. When a new input arrives, it becomes the data for a new node which is then linked back by a new directed edge to the current node. The new node then becomes the current node. In this way, the first four inputs would generate the structure shown in fig. 1(a), the second '*a*' being added when '*c*' is the current node. However, with a threshold distance of 2, this second '*a*' will be discovered to be a repetition, by searching up to the threshold distance from '*c*'. So rather than adding '*a*' as a new node, instead a link is created forward from the first '*a*', as shown in fig. 1(b), as though the second '*a*' node in fig. 1(a) had been picked up and placed on the first. Now with the first '*a*' as current node, we add '*d*' as shown in fig. 1(c) and then input '*c*' causes another identification to take place, resulting in fig. 1(d).

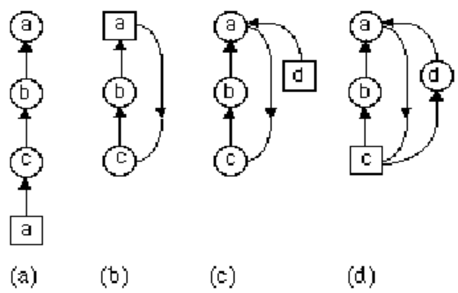

Fig. 1. Network built from input sequence $a\ b\ c\ a\ d\ c$, with threshold distance 2

By continuing the example we encounter some more idiosyncracies of net building. In fig. 2(a), '*e*' and '*a*' have been added. Once again there are two copies of '*a*', but this time they are too far apart for identification to have taken place. After a further input of '*c*', the second '*a*' gets a link forward from '*c*' (fig. 2(b)) and the effect is to bring the two copies of '*a*' within distance 2 of each other. However, no identification takes place because searching and identification must involve the current node and is only triggered by the addition of data at the current node. In fig. 2(c), the next input, '*e*', has been recognised as a repetition and the current node has accordingly moved to the '*e*' node; but no extra edge is added since the identification merely duplicates an existing edge. However, when '*c*' is now added, the identification, taking place in the opposite direction to this existing edge, does generate a new link (fig. 2(d)).

Now consider sending $e\ a$ to the network that we have built so far. The input '*e*' will move the current node to node '*e*' without adding a new edge. A search from this node for '*a*' will locate only the bottommost copy since the top '*a*' is beyond the threshold distance. Again, the current node moves but no edge is added. In fig. 3(a) the result is shown when a further two inputs, '*f*' and '*c*', have arrived. The next two inputs are '*a*' and '*c*'. On '*a*', a link is formed from the bottommost existing '*a*' (fig. 3(b)). There are now two copies of '*c*' within the

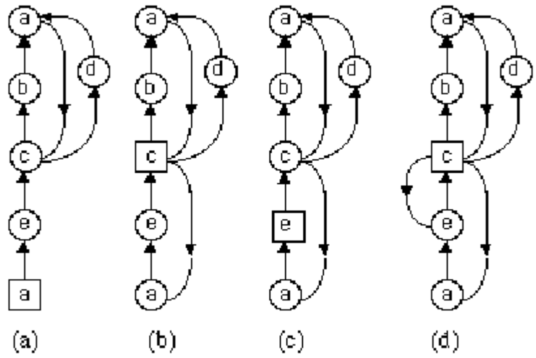

Fig. 2. Effect of presenting inputs $e\ a\ c\ e\ c$ to the network in fig. 1(d)

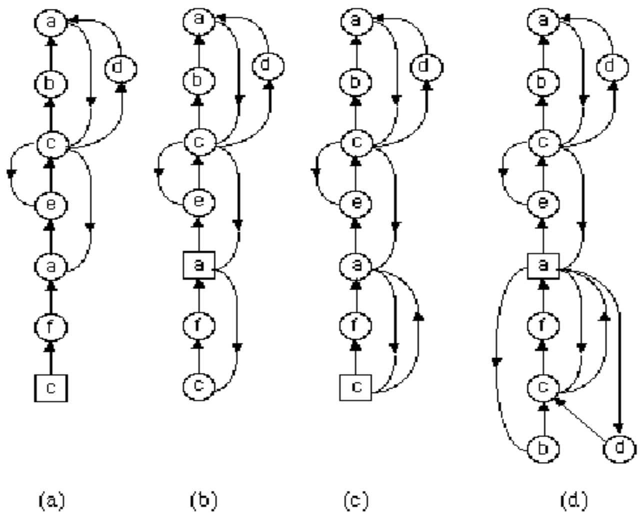

Fig. 3. Effect of presenting inputs $e\ a\ c\ e\ c$ to the network in fig. 2(d)

threshold distance of the current node. We will suppose that searching is carried out in a breadth-first manner, fanning out from the current node and terminating as soon as a repetition is found or the threshold distance is reached. According to this, input 'c' will produce the network in fig. 3(c): only an identification with the bottommost 'c' has occurred, since it is found earlier than the topmost 'c'. Finally, inputs $b\ a\ c\ d\ a$ produce the network in fig. 3(d).

We note that, given an input stream which is sufficiently repetitive, the simple principles with which we started can lead to quite intricate structures, with multiple copies of data within the threshold distance, nodes of arbitrary degree, cycles of any length, and so on.

We would contend that, given an evolutionary pressure to store inputs selectively, the forming of associative links to eliminate repetition is a plausible response[1]. There are two parameters: **network capacity** C (maximum number of nodes allowed in the network) and **node capacity** c (maximum number of edges which may be incident with a node). These constraints are crucial, since learning, to us, means reconstucting a worldview graph *incompletely*. What is left out of the reconstruction is as important as what goes in: including everything encountered during random walk would place an intolerable burden on the processes of cognition[2]. We must decide what to do when C or c are exceeded — for the former we ignore inputs requiring new nodes to be added; for the latter we delete the oldest link. These policies could also be selected by evolution.

3 Experiments: Learning in Association Networks

3.1 First Experiment: The "World Graph"

With our main assumption stated and the basic model described, we can wonder what its learning ability might be. Learning is seen as reading a training set presented as a stream of elements in temporal sequence, drawn from the original graph by generating a random walk along it. The AN is expected to replicate the graph as faithfully as possible, and its learning capacity will be scored from the resemblance between the constructed graph and the original one. We are not bound to concede the actual existence of such a structure according to which a particular aspect of the "natural world" would be organised, except from a rather speculative viewpoint. Our definition assumes that every learning process may be seen as that of inferring a highly dimensional tissue of concepts from a one-dimensional, temporal sample random walk of arbitrary, though finite length.

The first experiment was to generate a random graph with 500 nodes and a maximum, Δ, of 40 edges per node, up to a total of 10000 edges (half the number of nodes times the number of edges per node). Random walks were produced with a total length of 30,000 steps (nodes) and the parameters of the AN were C, c (see section 2) and the semantic threshold, defined as the minimum number of consecutive time units that a node has to be **accessible** (i.e. connected to the graph) in order to be deemed as a **semantic** (i.e. stable) item of memory.

Fig. 4 shows the performance of the AN as a function of time, for c at 5, 10, 20 and 30, with C=500 (same size as in the original graph) and semantic threshold at 3000. We can describe the dynamics as a cyclic, quite regular, alternation between periods of rapid growth of the accessible memory and simultaneous and slower growth of the semantic one, and sudden collapses leading to dramatic loss of memory. Two phenomena are easily observable: first, the evolution of semantic memory replicates that of accessible memory, but with a smaller amplitude; this

[1] A similar algorithm has been independently proposed for data compression by Mojdeh [8].

[2] We are reminded of the story "Funes, the Memorious" by Borges, in which Funes, who remembers everything, is completely unable to operate intelligently.

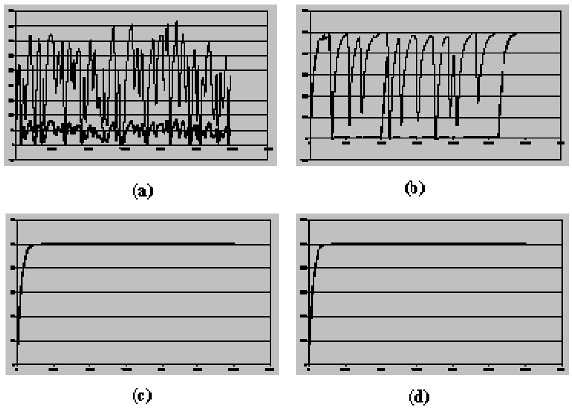

Fig. 4. Reconstructed graph from random walk on near-regular graph. Upper (light): accessible memory, lower (dark): semantic memory. Number of Nodes in graph: 500. Node Capacities: (a) 5, (b) 10, (c) 20, (d) 30.

is particularly apparent for low values of c. Second, the frequency of collapse decreases with increasing c, as one would expect, until some critical value beyond which the system acquires a core of stable memory: no matter whether any collapses occurred in the past, the system can be deemed to have reached an equilibrium point and the probability of a new collapse becomes negligible.

3.2 Second Experiment: A Case of "Real Life"

Our second experiment aimed at testing the ability of the AN to infer complex topological relations from a single input stream, and the practical problem chosen was that of the tourist having to construct a mental representation of the topology of a town (see the Introduction) from a set of bus routes presented as a temporal sequence of relevant bus stops drawn from those routes.

An input file was produced by concatenating all the public bus routes (red buses) serving London and Greater London. After appropriate filtering and standardising, the file was in condition to be processed so as to generate the "world graph", whose nodes are the stops and that has an edge connecting a pair of nodes if both are consecutively included in at least one route. This graph contained 430 nodes with degrees ranging from 2 to 10. Again, a random walk was generated from this graph, with a total length of 30,000 steps, as in the previous experiment; the parameters were also the same.

Fig. 5 shows the performance of the memory as a function of time, for c at 6, 7, 8, 9 and 10, with C=430 and semantic threshold at 3000. We can roughly observe the same general behaviour as in the previous case (randomly generated regular graph). However, whilst in the previous case a c of at most 50% of the maximum degree of the original graph ($c/\Delta = 20/40$) was enough for the system to eventually overcome the cyclic phase of alternating "memory growth - collapse", in this "real case" only by raising the node capacity to 80% of

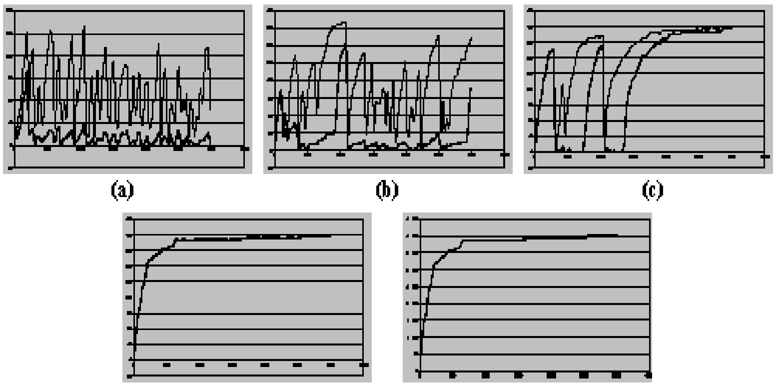

Fig. 5. Reconstructed graph from random walk on Bus Routes. Upper (light): accessible memory, lower (dark): semantic memory. Number of nodes in graph: 430. Node capacities: (a) 6, (b) 7, (c) 8, (d) 9, (e) 10.

the maximum degree (8/10) does the system succeed in reaching a stable regime. We explain this difference: in the random graph the variance of node degrees was small, the graphs being close to regular. Conversely, the variance for the "world graph" of London is so large that, given a certain ratio c/Δ, the probability of becoming disconnected and collapsing in a few steps is, on average, much higher.

We can conclude that no fixed dependence can be stated between the critical value for c (to ensure assymptotic stability) and the size and maximum degree of the original graph, but also the variance of degrees has to be considered.

3.3 Phase Transition Diagrams

On the basis of the previous results (and similar ones), we can distinguish qualitative phase states for the dynamics of the learning system, i.e. different landscapes for the temporal evolution of the process of memory formation. For c/Δ below some r (that basically depends on the structure of the original world graph), a state of rather regular alternation between memory growth and memory collapse takes place. For $c/\Delta > r$, an equilibrium is suddenly reached and the system rapidly evolves to complete reconstruction of the node structure, within the limitations of connectivity imposed by the ratio c/Δ at which the AN is set.

We can properly speak of *phase transition*, moreover a first order phase transition, since no smooth change takes place between the cyclic, unstable phase and the stable one in terms of frequency of collapses, but rather when a certain critical period is reached, the system enters the equilibrium phase. For the bus routes case, for example, this critical value is typically around 7000 steps.

Fig. 6 shows the diagrams. In the regular graph, the border line is more definite since we had more values of c/Δ. As for the London graph, the line is straight: the dynamics changes abruptly from transition point at t $\sim$ 13000 for $c/\Delta = 0.8$ to no transition (process stable from the beginning) for $c/\Delta = 0.9$.

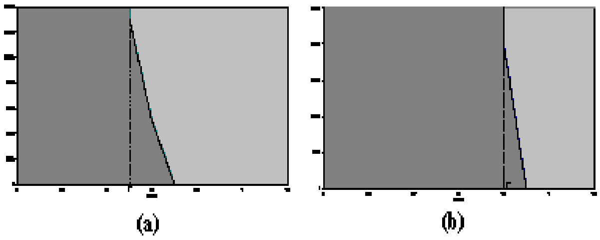

Fig. 6. Transition diagrams for the dynamics of memory formation: (a) Randomly generated regular graph (500 nodes) (b) "World Graph" for Central London (430 nodes)

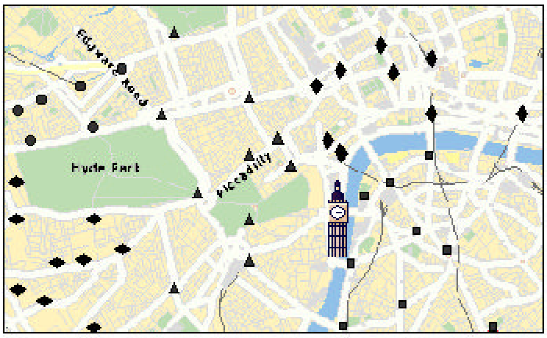

Fig. 7. Organisation of main places of London, according to neighbourhood relations, in five clusters, represented by: circles (top left), triangles (middle), vertical diamonds (top right), horizontal diamonds (bottom left), squares (bottom right)

3.4 Clustering and Classification

The AN can be used for clustering and classification. By virtue of the topological way in which it organises the information, the AN can extract relationships between concepts that are implicit in the temporal order of the input stream.

Two types of experiments were carried out: *clustering* and *association*. The first consisted of presenting an input sequence to the system and causing it to construct the corresponding AN. Then, standard clustering was applied, using as distance between two nodes the length of the shortest path between them. Fig. 7 shows the organisation of the main places of London in clusters, according to their neighbourhood relations, as derived from the bus routes network.

The second experiment, association or stimulus-response, was undertaken on the same network and consisted of entering a new input and searching for its immediate neighbours in the graph. A new parameter, called **semantic relevance** (SR), is used for rating the significance of an item of information: it is set initially to zero for every new item, and is decreased by 1 each time the same information appears in the input stream. If the SR has fallen below a fixed threshold, the replicated data is no longer added to the network. The effect of SR can be understood by using the AN in the stimulus-response (association) mode. With a threshold of SR of -10, the answer to *Westminster* was the set {LambethPalace, CabinetWarRooms, DowningSt, Archway, HorseGuards, VauxhallBridge, Victoria, HydeParkCorner}, while for a threshold of -5, the response was {BuckinghamPal, HydeParkCorner, LambethPalace, CabinetWarRooms, DowningSt}. Finally, for a threshold of 0 (only the first occurrence of each word is considered for constructing the net), the response was just {BuckinhamPal, HydeParkCorner}. Hence the SR prevents node degrees from growing too large, much like c: if the threshold is set to a value giving little or no restriction, e.g. -10, roughly the same three groups of answers shown for the question *Westminster* can be obtained with c at 10, 5/6 and 2, respectively.

Many fields are suitable for this model, e.g. in extracting relations between terms in a text. We took successive editions of the *Times* online to produce an input stream. After eliminating punctuation signs and changing capital to lower case, the input size was 7771 words. We used it to produce chains of associated concepts. For example, starting with the word "children" and ending in "judge", the sequence { *children - people - rights - laws - judge* } was obtained.

4 Discussion

We have introduced a system, the Association Network, that is able to model situations of learning from real environments, and to infer complex relations from a single input stream, as expected from an intelligent (human) agent. Our hypothesis is that, irrespective of the high parallelism of the neural structures involved in cognitive processes at the brain cortex, the channels through which information is conveyed from the environment to its final location (in whatever form of neural structure) can transmit only one data item per time unit.

Experiments showed the abilities of the AN: learning of complex topological structures; clustering and association; existence of a phase transition regime. This work is preliminary; behaviour in the limit could be analysed, given certain parameter settings. The WWW and other very-large-scale systems have given rise to new techniques in the field of random structures, ideal for this work [4][2].

Some comment is necessary on the somewhat metaphysical concept of "world graph", i.e. the implicit hypothesis that any domain of the natural world is isomorphic to some "mental model" and this is, in its turn, representable as a graph or a network of concepts. Maybe this concept is justifiable case by case; e.g., in the process of learning a language, some semantic network might be deemed as implicitly present in a dictionary and could be constructed from it.

We distinguish this model from that provided by associative or distributed memories, as being at a higher semantic level than these, or than reinforcement learning or evolutionary computation. Our approach is closer to schemata or frame models of memory, going back to Bartlett's work [1] and our dependence on association forming owes more to James [3] than to the PDP group [6].

That said, we must respond to the standard GOFAI (Good Old-Fashioned Artificial Intelligence) criticism that our model relies on the pre-existence of data concepts ('bus stop', 'football', 'chess piece') which themselves represent the core challenge for AI. In some cases, it might be appropriate to declare that the nodes of a certain AN were individual associative memories; the job of the AN being to provide a higher level interconnecting structure for these memories. Ultimately, the co-evolution of many communicating networks, at many levels of abstraction, performing different tasks might implement something like Minsky's Society of Mind [7], a (still) very persuasive view of how the human mind operates.

As for the emergence of different levels in an intercommunicating structure of AN's, we explored how semantic relevance (see subsection 3.4), rather than merely rejecting repeated items, might provide a filter, whereby such items are transferred as input to a higher level network (and from that network to one yet higher, etc). This might operate in processing language, where recurring words (e.g. articles or connectives) might be stored elsewhere both to avoid false associations between verbs and nouns and to capture their structural significance.

Our model is interactive: learning and application (cognition) are simultaneous, unlike the classical machine learning approach in which these two processes are consecutive. From the computational perspective, we propose a general way of structuring repetition rich sequential data. More important, but more speculative, is the possibility that we have invented a credible model of learning, i.e. concept formation in humans or, at least, given a major step in that direction.

References

[1] Bartlett, F.C.: Remembering: a Study in Experimental and Social Psychology. Cambridge University Press (1932)

[2] Bollobas, B. and Riordan, O., Coupling scale-free and classical random graphs. Submitted to Internet Mathematics (2003)

[3] James, W.: The Principles of Psychology, vol. 1. New York (1890), reprinted by Dover Publications (1950)

[4] Kleinberg, J.M, Kumar, R., Raghavan, R., Rajagopalan, S. and Tomkins, A.S.: The web as a graph: measurements, models, and methods. In Proc. 5th Int. Conf. on Computing and Combinatorics, Tokyo, Japan, 1999, Springer Verlag (1999), 1-17

[5] McClelland, J. L. and Rumelhart, D. E.: Distributed memory and the representation of general and specific information. J. Exp. Psych.: General, 114(2) (1985) 159-188

[6] McClelland, J. L., Rumelhart, D. E. and the PDP Res. Group: Parallel Distributed Processing: Explorations in the Microstructure of Cognition (2 vols). MIT Press (1986)

[7] Minsky, M.: The Society of Mind. Simon and Schuster, New York (1986)

[8] Mojdeh, D.: Codes and Graphs. 2nd European Workshop on Algebraic Graph Theory, Univ. Edinburgh (July 2001), (preprint, Univ. Mazandaran, Iran)