

ACCOUNTING FOR EPISODIC, SEMANTIC AND PROCEDURAL MEMORY IN THE RECOMMENDATION ARCHITECTURE COGNITIVE MODEL

L. ANDREW COWARD

Department of Computer Science, Australian National University, ACT 0200, Australia
andrew.coward@anu.edu.au

A qualitative explanation for experimental observations of episodic, semantic and procedural memory and the differences between them is provided based on the different ways in which information derived from the senses is recorded and accessed in the recommendation architecture cognitive model.

1. Introduction

The recommendation architecture (RA) is an information architecture which attempts to bridge the gap between psychological descriptions of cognitive phenomena and physiology (Coward, 1990). It represents a radically different approach from artificial neural networks, dynamical systems or computational modeling of symbolic systems. The different approach is needed because of the requirement to recognize and address the severe architectural constraints experienced by any system which learns large numbers of features with limited information handling resources. Information handling resources include resources needed for information recording, information manipulation (or “processing”) and internal connectivity. Theoretical arguments demonstrate that as the ratio of features to resources increases, the system will be constrained more and more tightly into the RA form (Coward, 2004a). If a cognitive model fails to take these constraints into account, it may be able to model one or a few features, or many features if there are no limits to available resources, but will break down when requirements similar to those on real biological brains are imposed.

There are strong resemblances between the physical forms of the RA and mammal brain physiology (Coward, 1990; 2001). Electronic implementations of the RA demonstrate bootstrapping of behaviour analogous with cognition with minimal a priori guidance (Coward, 2001).

This paper demonstrates that the phenomena and observed dissociations between semantic, episodic and procedural memory find a straightforward

explanation in terms of the available, electronically implemented information recording and access mechanisms in the RA cognitive model.

2. The Recommendation Architecture (RA)

In the RA there is a primary architectural separation between a modular hierarchy called clustering and a component hierarchy called competition. Clustering defines conditions within the information available to the brain and detects the occurrence of any defined condition. Modules in clustering detect sets of similar conditions, where the similarity between conditions within a higher module is somewhat less than within the more detailed modules that make up the higher level module. A subset of the conditions detected at any point in time is communicated to competition. Different components within competition correspond with general behaviour types, more specific types, and individual behaviours and sequences of behaviours. Any one condition detection by clustering is communicated to a wide range of components, where it is interpreted as a recommendation in favour of the corresponding behaviours, each recommendation having a different weight. Competition adds the weights of each recommended behaviour across all currently detected conditions, and implements the most strongly recommended behaviours. Consequence feedback following a behaviour can change the recommendation weights of recently active conditions into recent behaviours, but cannot change the definitions of the conditions.

Devices in clustering learn and respond in radically different ways from conventional connectionist device algorithms. A clustering device permanently records a set of conditions. In other words, once a condition has been recorded, an exact repetition of that condition will always be detected. Systems based on this type of device algorithm have been implemented electronically and demonstrated cognitive like capabilities (Coward, 2001). To be recorded, a condition must actually occur within the information available to the device, be sufficiently similar to conditions already recorded on the device, and the device must also be receiving signals encouraging it to record conditions. These signals come from special purpose devices indicating the general level of condition detection in other modules. A device is activated by the recording of and by any subsequent repetition of a condition. The output of an activated device is a series of spikes. The average rate of spike production indicates the number of conditions which are currently present, and a frequency modulation of the spike rate (i.e. bunching of spikes towards the peaks in a modulation frequency) indicates the input population within which the condition was detected. This

modulation is determined by the modulation of the inputs to the device. For example, if modulation is imposed only on the subset of the outputs from a visual input domain which correspond with an area within a closed boundary (i.e. a visual object), only conditions within the object will be detected by devices processing outputs from the whole domain, and recommended behaviours will be in response to this object. Conditions detected by the same physical set of devices within two different objects can be kept separate if different phases of frequency modulation are imposed on inputs from the two objects (Coward, 2004b).

Careful management of when and where conditions will be recorded is essential. Additions are driven by the requirement that an adequate range of behavioural recommendations be available in response to every input state, which is equivalent to a requirement that the number of conditions detected in response to every input state must be above a minimum level. Additional conditions are recorded until this level is reached. Recording of conditions takes place in the modules containing the most similar conditions. New modules are created if conditions dissimilar to those in any existing module conditions are required. The location of condition recording is determined by the relative activity in a range of alternative locations (Coward 2001).

Devices in clustering are arranged in layers in which the first layer receives sensory inputs, and condition defining inputs to each subsequent layer come from just the preceding layer. All conditions detected within one layer are within the same range of complexity, where the complexity of a condition is the number of sensory inputs (including duplicates) that contribute to the condition either directly or via intermediate conditions. Parallel arrays of columns extend through several layers. Columns manage when and where additional conditions will be recorded, and in addition they detect and communicate the detection of conditions within a specific range of complexity to competition. These conditions are called the portfolio of the column, and arrays add conditions until some portfolios are detected within every input state.

These device algorithms and change management processes have all been implemented in electronic form (Coward, 2001; 2004b).

3. Definition of Conditions and Portfolios

A vast range of inputs containing information about the external environment and the internal state of both the brain and the body are available to the brain from the senses etc. A somewhat oversimplified way of understanding the definition of conditions is that one condition corresponds with each input in

a specific set of inputs being present to an individually specified degree. Because conditions cannot be specified a priori, there is a random element to the definition of conditions, and because conditions are not changed after being recorded, any one condition or portfolio is perceptually, cognitively and behaviourally ambiguous. Unambiguous meanings are only achieved in competition, using populations of portfolios. However, conditions with complexities of the same order of magnitude as visual object perceptions will tend to be less ambiguous with respect to categorization of visual objects than conditions on other levels of complexity. Although no conditions on any level correlate unambiguously with such categories, conditions on such levels of complexity will provide better discrimination between such categories.

This simple view of conditions is made much more complex because a column can be activated not only by the presence of its conditions within sensory inputs, but also indirectly by two types of mechanism. One mechanism is activation if a number of other columns are already active which were often active in the past at the same time as the column. The other mechanism is activation if a number of other columns are already active which recorded conditions at the same time in the past as the column. An available variation is if the past activity or condition recording of one portfolio was shortly before or after that in the other. Such indirect activations are behaviours which must be recommended by already active columns to competition and accepted.

When there is simultaneous activity or condition recording in two columns, there is a strong recommendation weight created initially in favour of the indirect activation, but such weights decline fairly rapidly with time. However, if an indirect activation occurs shortly before a behaviour which is followed by positive consequences, the decline is reduced. Frequent such occurrences stabilize or increase the weight.

These indirect activation mechanisms can be viewed as supplementing the conditions present in current sensory inputs with other conditions which have a significant probability of being relevant to determining the most appropriate current behaviour.

A newly recorded condition is made up of a set of currently active component conditions. Some of these component conditions may be combinations of currently present sensory inputs, and others could have been activated by one of the two indirect mechanisms. Both the definition of conditions in terms of sensory inputs and the relationship between sensory inputs and the resultant pattern of condition activation can therefore be very complex.

4. The RA cognitive model

In the RA, information can be therefore accessed by four qualitatively different mechanisms. Firstly, the actual presence of a condition within current sensory inputs activates portfolios containing the condition. Secondly, an activated portfolio can recommend activation of other portfolios which were often been active at the same time in the past. Thirdly, an activated portfolio can recommend activation of other portfolios which recorded conditions at the same time in the past. A variant of these mechanisms is activation on the basis of activity by a portfolio somewhat before or somewhat after activity by the other portfolio. Indirectly activated portfolios can in turn recommend activation of yet other portfolios. The fourth mechanism is comparison of recommendation weights. The weights of all active behaviours into each recommended behaviour are totaled, and the behaviours with the strongest weights are implemented.

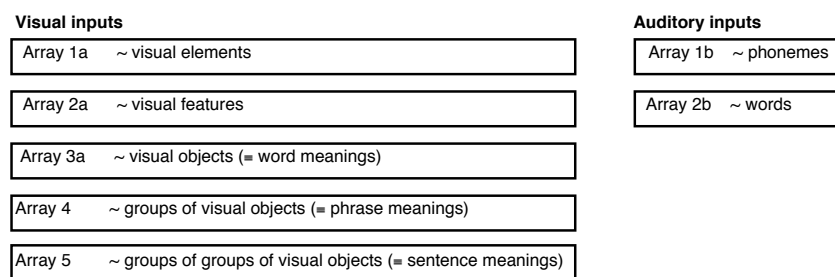


Figure 1. Architecture to support cognitive processes. Arrays detect conditions on different levels of complexity, with condition defining information passing sequentially from top to bottom. The outputs of an array indicate the detection of conditions within roughly the same range of complexity as the indicated cognitive category (features, objects, groups of objects etc.) but conditions do not correlate unambiguously with such categories. Conditions are detected independently in the visual and auditory domains.

The simplest arrangement of clustering arrays to support cognitive behaviour is illustrated in figure 1, and examples of competition subsystems associated with one array are illustrated in figure 2. In figure 1, different array outputs indicate the detection of portfolios with a complexity able to discriminate between cognitive features, objects, groups of objects etc. but individual portfolios do not correlate unambiguously with such objects etc. Portfolios able to discriminate between different groups of objects record conditions containing information derived from each member of groups.

Behaviour types which could be recommended by a portfolio through connectivity to the corresponding component include: prolong the activity of some currently active portfolios; activate portfolios active at the same time in the past as the currently active portfolio; activate portfolios containing conditions recorded at the same time in the past as some conditions in the currently active portfolio; activate portfolios containing conditions recorded just before or just after some conditions in the currently active portfolio; synchronize the activity (i.e. phase of modulation frequency) of several different groups of currently active portfolios; release outputs from one array to the next array; perform a general sequence of attention behaviours; perform a specific sequences of attention behaviours; perform an individual attention behaviour; speak a word; and say a phrase.

Competition receives outputs from portfolios currently being detected by clustering. If a portfolio has been present in the past at the same time as the performance of a number of different behaviours, it will have acquired recommendation weights in favour of or against those behaviours in the component corresponding with the behavior, depending on the consequence feedback from those behaviours. Competition adds the weights of all currently recommended behaviours and selects the behaviour with the largest weight. New portfolios are given an initial weight similar to the weights of the most similar previously existing portfolios, or genetically defined initial weights. If portfolios are new and no significant recommendation strength has been assigned in these ways, target behaviours are selected randomly. Such random selection can be limited to behaviours within a behaviour type which has already been selected. Alternatively, a behaviour can be selected by imitation of an externally observed behaviour.

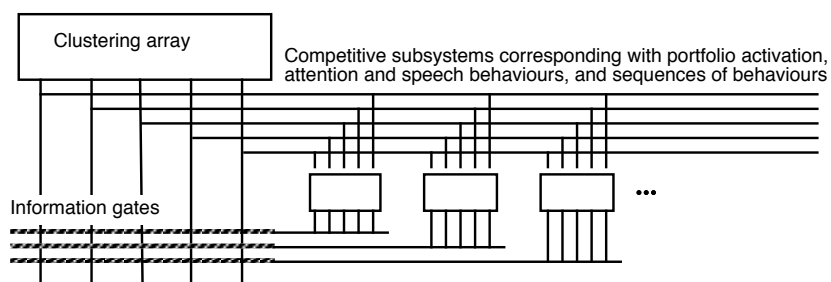


Figure 2 Competitive components receiving outputs from one array. Different behavioural interpretations are placed upon the same clustering outputs by different components. Competitive inhibition between and within competitive components limits selected behaviours to a small,

consistent set. The behaviour accepted by a competitive subsystem is implemented by release of the outputs from clustering which correspond with the behaviour to a clustering level or to motor behaviour. This release behaviour is indicated by the information gates.

5. Different Types of Memory

The permanent recording of conditions means that the system has the capability to learn from single experiences. The model implies that there will be a level of condition recording in response to every experience, with higher levels for experiences with higher degrees of novelty. The level of condition recording therefore correlates with the degree of novelty of the experience, accounting for the high human capability to detect the novelty or otherwise of an experience (Standing, Conexio & Haber, 1970).

5.1. *Episodic Memory*

Episodic memory is memory of the past with a context of what else happened at the time, in contrast with semantic memory in which memories of facts are detached from memories of where those facts were learned (Tulving, 2002). Various types of recall experiments measure episodic memory.

In targeted recall, subjects are asked to recall particular past events (Neisser & Libby, 2000). In cued recall, subjects are given a cue which may be a word (Robinson, 1976) or a type such as "vivid memories" (Rubin & Kozin, 1984). In involuntary recall, some specific environmental stimulus such as a smell or a taste brings a memory to mind unsought (Berntsen, 1996).

The starting point for targeted recall in the RA model is hearing words which describe an event. Portfolios are activated which contain conditions within the sounds of the words. Secondary populations of portfolios containing conditions which occurred within visual and other sensory inputs are activated on the basis that the secondary portfolios have often been active in the past at the same time as the primary "auditory" portfolios. A significant proportion of the portfolios in these secondary populations were also active during the target event, and a somewhat smaller proportion recorded conditions during that event. Because of the words used, the proportions are larger for the target event than for any other event.

All active portfolios have recommendation strengths in favour of activating other portfolios which recorded conditions at similar times in the past. Active populations at higher levels derived from the presence of words like "recall" have recommendation strength in favour of accepting these types of recommendations. Because the target event has the highest proportion of

activated portfolios in the secondary population, acceptance of such recommendations will tend to result in an active tertiary population with an even higher proportion of portfolios which recorded conditions during the target event. This process is self reinforcing, especially if a large number of conditions were recorded during the target event. The resultant population will be experienced as a re-perception of the original event, although the portfolios closest to sensory input are not reactivated, and there is therefore no visual hallucination of the event. The activated portfolios in this population have recommendation strength in favour of, for example, generating verbal descriptions of the event. Use of recommendation strength in favour of activating portfolios which recorded conditions somewhat before or somewhat after condition recording in currently active portfolios allows the re-perception to be set at the beginning of the event and to be moved through the event.

Recommendation strengths will always also be present in favour of activating portfolios on the basis of simultaneous past activity and simultaneous past recording during other events. The activated population is therefore unlikely to be an exact match for the original, although in general the higher the level of condition recording during the event, the greater the probability of a close match.

Cued recall operates in a very similar fashion. However, the initial secondary population may contain portfolios which recorded conditions during a number of past events. In the absence of specific indication of one event in the verbal cue, the tertiary population will evolve towards the event which happened to be represented by the highest degree of condition recording in the initial population. Events which resulted in a high degree of condition recording across many portfolios will tend to be the end points of this process.

Involuntary recall is the result of strong condition recording in portfolios activated in response to a sensory stimulus (for example, a novel smell or taste) at the same time as strong condition recording in other portfolios in response to some event. A later repetition of the sensory stimulus activates the portfolios which originally responded to that stimulus. These portfolios in turn activate portfolios which recorded conditions at the same time in the past, resulting in a re-perception of the event.

5.2.Semantic Memory

A typical way of measuring semantic memory in the laboratory is sentence and category verification. Sentence verification experiments measure the time for subjects to respond with the correctness of sentences like "Is a robin a bird?"

or “Is a penguin a bird”. Category verification experiments are essentially equivalent and measure such times for simple category-exemplar pairs like bird-robin or bird-tree. It is found that for different members of the same category paired with the correct category, responses are faster for more typical category exemplars. For example, the response to bird-robin is faster than to bird-chicken. However, responses to clearly incorrect category-exemplar pairs like “Is canary an animal” is also fast (Rips, Shoben & Smith, 1973).

The RA model for category verification experiments can be understood by considering the portfolio populations activated in response to the words indicating category and exemplar. The portfolios activated in response to hearing the name of the category are portfolios which have often been active in the past when the category name was also present. Hence they will be portfolios also present when exemplars of the category were present, since this is the way the category is learned. Hence the active population is the set of portfolios which have most often been present when different category exemplars have been present. There will therefore be an overlap between the population activated in response to the name of the category and that activated in response to the name of the exemplar. This overlap will be greater for more typical exemplars, and very small if the exemplar is not a member of the category. The degree of overlap is itself a condition which can be detected and recommends for or against identifying the exemplar as a category member. If the exemplar is typical, overlap is substantial and its detection is rapid. Atypical exemplars have more moderate overlap, and more time may be required to expand the portfolio populations to include portfolios active at the same time in the past but slightly less often to achieve an overlap adequate to generate the appropriate verbal response. However, objects which are not in any way members of the category will have negligible overlap which again is detected rapidly. The model is therefore in agreement with the observations of Rips, Shoben & Smith (1973).

In contrast with the spreading activation model of Collins & Loftus (1975), there are no units which correspond with concepts like categories or the features of categories. Portfolios are groups of conditions in which there has been a degree of randomness in the definition of individual conditions, but conditions within one portfolio have some similarity with each other and have tended to occur at similar times in the past. A portfolio may therefore have a probabilistic correlation with many different features and categories, with the probabilities expressed, for example, as recommendation weights into naming the features or categories.

5.3.Procedural Memory

Procedural memory is defined as the ability to acquire skills. Observations of amnesics indicate that such memory is at least partially dissociated from semantic memory, since amnesics can acquire such skills at apparently normal rates. Thus amnesics can acquire motor skills such as mirror tracing tasks (Corkin, 1968).

In the RA model for learning a skill, portfolios activated within clustering in environments where the skill is relevant must acquire weights in competition associated with skilled behaviours. These portfolios will generally include both new information elements resulting from novelty in the environments and information elements recorded in prior experiences. Although the new elements may be particularly useful for recommending the new behaviours, some skill learning would be possible using only previously recorded elements which happen to occur in the new environments. Thus skill acquisition could proceed in the absence of condition recording.

6. Conclusions

The RA provides a qualitative account for the phenomena of episodic, semantic and procedural memory and the differences between them. This capability results from the different mechanisms by which information derived from sensory inputs is recorded and accessed, with different combinations of mechanisms corresponding with the different memory systems.

References

- Berntsen, D. (1996). Involuntary autobiographical memories. *Applied Cognitive Psychology* 10, 435 - 454.
- Collins, A. & Loftus, E. (1975). A spreading-activation theory of semantic processing. *Psychological Review* 82(6), 407 – 428.
- Corkin, S. (1968). Acquisition of motor skill after bilateral medial temporal-lobe excision. *Neuropsychologia* 6, 225 - 264.
- Coward, L. A. (1990). *Pattern Thinking*. New York: Praeger.
- Coward, L. A. (2001). The RA: lessons from the design of large scale electronic systems for cognitive science. *Journal of Cognitive Systems Research*, 2, 2, 111-156.
- Coward, L. A. (2004a). The RA Model for Human Cognition. Proceedings of the Conference on Brain Inspired Cognitive Systems, University of Stirling, Scotland.
- Coward, L. A. (2004b). Simulation of a Proposed Binding Model. Proceedings of Workshop on Information Coding in Early Sensory Stages, BICS 2004, University of Stirling, Scotland.
- Neisser, U., & Libby, L. K. (2000). Remembering life experiences. In E. Tulving and F. I. M. Craik (Eds.). *The Oxford Handbook of Memory*. Oxford: Oxford University Press.
- Rips, L., Shoben, J. & Smith, E. (1973). Semantic Distance and Verification of Semantic Relations. *Journal of Verbal Learning and Verbal Behaviour* 12, 1 – 20.

- Robinson, J. A. (1976). Sampling autobiographical memory. *Cognitive Psychology* 8, 578 - 595.
- Rubin, D. C. & Kozin, M. (1984). Vivid memories. *Cognition* 16, 81 - 95.
- Standing, L., Conexio, J., & Haber, R.N. (1970). Perception and memory for pictures: single-trial learning of 2500 visual stimuli. *Psychonomic Science* 19(2), 73 - 74.