

A MODEL FOR REPRESENTATION OF CONCEPTS IN THE BRAIN

L. Andrew Coward¹ and Tamás (Tom) D. Gedeon²

¹Department of Computer Science, Australian National University,
Canberra, ACT 0200, Australia, andrew.coward@anu.edu.au

²Department of Computer Science, Australian National University,
Canberra, ACT 0200, Australia, tom.gedeon@anu.edu.au

ABSTRACT

Categorization is a primary application for conceptual knowledge stored in the human brain. Categorization is often based on similarity, involving apparent use of both prototypes and stored exemplars. Some human categorization appears to be rule or theory based rather than based on similarity. Attempts to model categorization often involve multiple subsystems to support the different observed approaches. Any system which must perform a large number of different but interdependent behaviours with limited information handling resources will tend to be constrained within a form called the recommendation architecture. Physiological structures in the brain resemble the forms of this recommendation architecture. The information recording and access mechanisms of the recommendation architecture result in all the different categorization phenomena and use only a single knowledge representation system. Categorization phenomena differ only in the ways in which information is accessed from the representation system.

1. INTRODUCTION

In brains, information derived from sensory inputs must be organized by physiological processes in such a way that higher cognition can be supported. This organizational structure is the way in which knowledge is represented. For any system, whether electronic or biological, solving a particular type of problem is made easier by an appropriate choice of knowledge representation. A typical illustration of this point is that simple additions are easier with Roman numerals, long division is much easier using Arabic numerals than with Roman numerals, and binary numbers are the best way to represent numbers in a computer. However, there is no representation which is ideal for every type of problem.

One major issue for the human brain is therefore how to organize information derived from sensory experiences in a way which is adequate to support the very wide range of different problems which the brain must solve. Some of these problems which must be solved are the needs to generate appropriate physical behaviours while moving about the world, acquiring food, avoiding threats, and engaging in complex social

behaviours etc. Others include higher cognitive processes such as categorization, speech processing, and reasoning. Equivalent information could of course be organized and stored in many different ways to support different types of problem, but this approach will tend to be ruled out by resource limitations.

A further issue is the need to define the content of the knowledge representations heuristically. This issue strongly interacts with the resource issue. The content of the knowledge representation in the brain must change to support learning, but because resource limitations mean that the same representation must be used for solving many different types of problem, changes which benefit one type of problem can easily result in undesirable side effects on the capability to solve other problem types.

The tension between resource limitations and need for learning without side effects forces any system which must learn many different behaviours within a set of architectural constraints called the recommendation architecture [1]. Electronic systems with the recommendation architecture have been implemented in software and their capabilities confirmed [2; 3; 4]. The brain must learn many different types of behaviour, and natural selection will tend to favour brains which require fewer resources. The human brain will therefore tend to be constrained within the recommendation bounds and there are in fact some strong resemblances between brain physiology and recommendation architecture forms [3; 5].

There have been a range of categorization procedures observed in human subjects, including evaluating whether an object resembles a prototype defining the category, resembles a range of different recorded exemplars of the category, meets rules defining a category, or has features which can be explained by a theory defining the category. It has been proposed that these procedures require different brain systems. Such a multiple subsystem approach raises the issue of whether the different subsystems can make use of the same information structures, or if more and/or duplicate information must be recorded to support the different subsystems.

The objective of this paper is to demonstrate that the limited range of information recording and access mechanisms utilized in the recommendation architecture model support all the different categorization pro-

cedures using a common knowledge representation structure. The differences between the procedures reflect different ways in which the same information is accessed. The observed categorization phenomena thus result naturally in any system within the recommendation architecture bounds.

2. THE THEORY OF CONCEPTS

The human mind is capable of thinking about an immense range of concepts. Some concepts are everyday (such as dog or house). Others can be distant in space or time (such as stars or extinct animal species). Yet others can be intricate abstractions from everyday experiences (such as democracy or the number pi), and others can be non-existent but imaginable (such as elephants with wings).

Attempts to create a theory of how concepts are defined began in antiquity. In the classical theory, concepts are complex mental representations with a structure that specifies a group of necessary and sufficient core characteristics for the concept to be applicable. These characteristics may themselves be concepts made up of more detailed characteristics. For example, the concept *BACHELOR* would be a complex mental representation with characteristics *UNMARRIED* and *MAN*. When categorizing a perception, a concept is accessed and decomposed into its characteristics, and the perception checked for the presence of all these characteristics. If all are present, the perception is categorized as an instance of the concept. A concept is learned by bringing together a group of pre-existing concepts to define a new concept. The major problem with this classical theory is the inability to define indisputable core characteristics for any real categories. A key example of this is the category *GAME* as discussed by Wittgenstein [6]. Wittgenstein considered a number of plausible definitions in terms of core characteristics, and demonstrated that in every case a clear example could be found of a game which did not meet the definition.

Psychological observations also provide difficulties for classical concept theory. Some of the key issues are that subjects identify typical members of a category more rapidly than less typical and subjects tend to produce more typical category instances when asked for examples [e.g. 7].

Prototype theory [8] attempts to overcome some of the difficulties with the classical theory. In this theory, concepts are defined by prototypes made up of a list of features, each feature having a weight. A perception would be categorized as an instance of the concept if and only if it possessed a sufficient number of these features, weighted for their importance. In one version of prototype theory a concept such as *FRUIT* would include attributes like *contains seeds*, *is sweet*, *grows on trees* and *is round*. An instance would be categorized as fruit if and only if it possessed a sufficient number of these characteristics weighted for their importance [9].

Thus category membership of a concept is determined by computing a measure of similarity based on the degree of feature match and comparing the degree of match with a threshold band of values. A similarity above the band indicates category membership. A similarity below the band indicates non membership. A similarity within the band indicates that membership depends upon context or occasion. Prototype theory can therefore provide an account for the inherent fuzziness of many categories, as in the observations of McCloskey and Glucksberg [10] who showed that subjects disagree on categorization, but generally only for less typical instances.

In exemplar theory, concepts are defined by the recording of a range of instances of the concept. Objects are classified on the basis of their similarity to these stored exemplars [11]. A key advantage of exemplar theory over prototype theory is that it provides an account for the observations that in human beings the actual instances perceived in a learning phase have a strong effect on categorization accuracy during a subsequent test phase. For artificial category learning in the laboratory [e.g. 12], new instances of a category which are strongly similar to individual training instances but less similar to the category average are identified more accurately than new instances which are more similar to the average but strongly similar to fewer individual training instances. Thus in these experiments, similarity to specific examples is more important than similarity to the category prototype.

A similar effect has been observed in real experience. In the experiments of Brooks, Norman and Allen [13], a number of photographs of different dermatological conditions with known diagnosis were obtained. Groups of four photographs were created, all instances of the same condition. Within a group there were two pairs of photographs. The photographs within a pair were similar, but the pairs were more different. In a training phase, doctors with dermatological experience were exposed to one photograph from a number of different groups and told the correct diagnosis. Later they were shown different photographs from the same groups as in the training phase and asked to select a diagnosis. Performance was better if the test photograph was more similar to the training photograph.

However, as pointed out by Brooks, Norman and Allen, a prototype model in which recent instances resulted in disproportionate changes to the weights of prototype features could also account for the observations.

A number of workers have also claimed a closer quantitative experimental match to predicted performance profiles for exemplar theories in laboratory experiments, but Minda and Smith [14] have argued that if prototype and exemplar models with the same parameters and power are used, this difference disappears.

Some recent theorizing on similarity based perceptual classification has generated a number of mod-

els with multiple categorization systems [11]. These models tend to have one subsystem which categorizes on the basis of rules or prototypes, and a second subsystem which uses more specific representations like stored exemplars. The problem with such models is that they are so flexible, with numerous adjustable parameters, that they can be made to match any experimental result.

The classical, prototype and exemplar theories all conceive of concepts as being based on similarity, and concept learning and use is based upon determining the similarity between a stored representation of a category and the particular instance to be categorized. However, this does not always appear to be an adequate basis. Thus Barsalou [15] pointed out that there are what he called ad hoc categories such as THINGS TO TAKE ON A CAMPING TRIP. There may be minimal apparent similarity between objects in the same conceptual category. Furthermore, some categories are defined by rules rather than similarity. For example, Tienson [16] pointed out that a triangle with a small part of one corner cut off by a line is categorized as a quadrilateral although it still strongly resembles a triangle. Keil [17] pointed out that people have theories that embody relationships between properties and determine categorizations. For example, when presented with a dolphin or whale with the appearance of a fish but the insides and lineage of a mammal, adults categorize it as a mammal.

Such dissociations between similarity and categorization imply the need for a rule based or theory based models as well as similarity based models. Models for mature concepts therefore tend to have two components [18]. There has been a tendency for similarity based information to be viewed as primary, based on the assumptions that similarity based categorization is learned first by children and similarity based information is accessed faster. However, Keil, Smith, Simons and Lev in [19] present both theoretical and experimental reasons why neither of these assumptions are correct.

There has thus been a tendency for models of concepts and the categorization process to become very complex, with multiple, semi-independent subsystems required to provide an account for psychological observations. Thus Smith, Patalano and Jonides [20] suggest that there are four procedures for determining whether a test object belongs to a particular category. These are:

- Determine the similarity of the object to the category prototype.
- Determine the similarity of the object to remembered category exemplars.
- Determine whether the object fits a rule defining the category.
- Determine whether the features of the object are best explained by the “theory” that underlies the category.

This modeling approach raises the question of whether these procedures can share information derived from sensory inputs.

3. THE RECOMMENDATION ARCHITECTURE

The recommendation architecture is a set of architectural bounds within which any system which must learn a complex combination of behaviours with limited resources will be constrained [1]. In other words, as the number of different behaviours increases, unless there are enough resources to allow behaviours to operate independently with duplication of similar information storage and processing for different features, the recommendation architecture becomes the only way to avoid catastrophic interference between early and later learning. A detailed description of a software implementation of a system with the recommendation architecture is available in [3].

At the highest level, a system with the recommendation architecture separates into two major subsystems called clustering and competition. Clustering selects, records and detects information conditions within the input space available to the system. A condition is defined by a subset of the sensory inputs available to the system, each input having a specified state. The condition occurs if all the inputs are in their specified state.

Many of these conditions are used to determine when and where additional conditions will be recorded, and a small subset of condition detections are released to competition. Competition interprets each condition detection as a recommendation in favour of a range of different behaviours, each with a different weight. Competition identifies the behaviour with the strongest total recommendation weight across all currently detected conditions and implements that behaviour. Recommendation weights are changed by consequence feedback following a behaviour, but conditions are not changed by such consequence feedback.

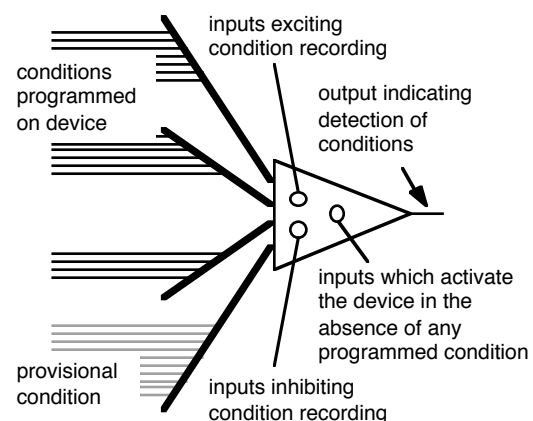


Figure 1. A condition recording device

As discussed in detail in [1], once a condition is recorded it cannot be changed. In other words, if a condition has been detected once, a future exact repe-

tion will always also be detected. The reason for this condition permanence is that when a condition is recorded it may acquire a range of different behavioural meanings in competition. Subsequent changes, for example to improve one behaviour, will in general result in undesirable side effects on other behaviours which outweigh any benefit. Close management of condition recording is therefore required to ensure that recording is limited and conditions are as behaviourally useful as possible.

This permanence of condition recording means that clustering is similarity based.

The number of possible conditions in a large input space is enormous. If conditions and behaviours are defined heuristically, there will be a random element to the selection of conditions. Because conditions are not subsequently changed, they will not correlate exactly with cognitive features or circumstances in which a particular behaviour is appropriate. All conditions are therefore cognitively and behaviourally ambiguous, and unambiguous meanings are only achieved by interpretation of groups of conditions in competition.

Within clustering, a device records a set of similar conditions. In this context two conditions are similar if a high proportion of their inputs are the same or tend to occur frequently at the same time. Such a device is illustrated in figure 1. Different groups of inputs define different conditions. In addition, somewhat larger groups of inputs define provisional conditions which will not be detected unless some special circumstances occur. The device has change management inputs that excite the recording of additional conditions, and other change management inputs that inhibit such recording. These change management inputs are derived from different groups of condition recording devices and indicate the average level of activity within the group.

If a high proportion of the inputs to a provisional condition are active, if no regular conditions are being detected, if inputs exciting condition recording are active and if inputs inhibiting such recording are inactive, a new condition will be recorded. This condition is the subset of provisional inputs which are active, and any subsequent repetition of the condition will be detected independent of the status of the change management inputs. A device also has inputs which can activate it in the absence of any of its programmed behaviours. Such activations are behaviours which must be accepted by competition. As discussed below, such activation behaviours expand the range of conditions available to influence cognitive behaviour on the basis of past activity of the conditions at the same time as currently active conditions.

Condition detecting devices are organized in layers. One layer detects conditions which are combinations of the conditions detected by the preceding layer. Layers therefore detect conditions on different levels of complexity relative to sensory inputs. A hierarchy of modules is overlaid on the layer structure. A small area of one layer forms a layer module. A sequence of

layer modules forms a column module. A parallel set of column modules forms an array module, as shown in figure 2. A sequence of arrays detecting conditions within different ranges of complexity form an area module.

A specific layer module produces outputs to competition, the set of conditions detected by that layer module are called the portfolio of the column. Such an output indicates that one or more portfolio conditions are present, but does not identify the conditions in detail. Conditions within one portfolio are similar and/or tend to occur at the same time, and tend to be different from the conditions in different portfolios. Conditions are recorded in response to a sensory input state if no existing portfolio is detecting conditions, and are recorded in the portfolio which contains conditions most similar to conditions actually occurring in the input state. If no conditions in the input state are sufficiently similar, a new portfolio can be initiated. Various layer modules within a column excite condition recording in their own column or inhibit condition recording in other columns or their own column in such a way that the similarities within a column portfolio and differences between portfolios are achieved, as described in detail in [3].




Figure 2. A group of column modules across a sequence of four layers

There are some conceptual similarities between portfolios and independent components analysis [21] in the sense that portfolios decompose a sequence of input states into partially statistically independent “features” in an unsupervised manner. The critical differences are firstly that portfolios can constantly evolve by addition of new conditions. This evolution means that new types of input states can be decomposed using an existing portfolio framework. An array of columns evolves in such a way that a number of portfolios are detected in every input state. The second difference is that portfolio evolution occurs in such a way that the occurrence of any previous portfolio definition will be detected by the latest portfolio. In other words, the volume of the input phase space to which a portfolio responds can expand but not contract. Behavioural interpretations placed upon the detection of a portfolio are therefore relatively stable.

A third difference is that although portfolios are relatively stable, additional portfolios can be defined in the similarity space of a group of portfolios if there

is an indication of need. Such an indication would be if the group was often active prior to one particular behaviour, and in these circumstances the consequences of the behaviour were sometimes positive and sometimes negative. The implication is that the existing portfolios do not adequately discriminate between behaviourally significant input differences. Adding portfolios does not affect the use of the existing portfolios for other behaviours.

Portfolios will not correspond exactly with, for example, cognitive features. No one portfolio will be present if and only if a particular feature is present. Rather, there will be a group of portfolios which tend to be present more often when a particular feature is present. Weights in favour of behaviours appropriate to the presence of the feature can be given to all such portfolios, and high integrity behaviour management achieved [3].

For example, portfolios defined in response to experiences of different instances of different types of FRUIT would not correspond exactly with characteristics like sweet, round, contains seeds, red etc. Rather, a portfolio would be initiated around an information condition which happened to occur in a sensory experience (which might or might not correspond with a fruit) and would be expanded by addition of conditions which were similar and/or happened to occur at the same time. One portfolio might contain some conditions correlating with redness, some with roundness. Such a portfolio would be activated in response to a red apple, but also a red ball. Such a portfolio would acquire recommendation strengths in favour of behaviours appropriate to both apples and balls. Appropriate behaviour is determined by the total recommendation strengths across all currently active portfolios. However, although portfolios do not correlate with features, they must be defined in such a way that although they can discriminate between any two sensory experiences with different behavioural implications, they not proliferate excessively. The structures and processes of the recommendation architecture ensure that this compromise is achieved [3].

The stability of portfolios under change means that activation of portfolios on the basis of past activity at the same time as currently active portfolios can be meaningful and behaviourally useful. These activations are themselves behaviours which must be recommended by an adequate population of currently active portfolios. Any one portfolio will acquire recommendation strengths in favour of a range of different behaviours, and in this sense any one portfolio is behaviourally ambiguous.

4. INFORMATION RECORDING AND ACCESS MECHANISMS IN THE RECOMMENDATION ARCHITECTURE

If a condition within a portfolio occurs in an input state, the portfolio will be activated. This portfolio will have recommendation strengths in favour of be-

haviours which were appropriate in response to similar input states in the past. However, there may be other portfolios which could provide relevant recommendation strengths in response to the current input state which are not directly activated by condition occurrence. Such other portfolios include portfolios which have recently been active at the same time as an active portfolio, portfolios which have often been active in the past at the same time as a currently active portfolio, and portfolios which have recorded conditions in the past at the same time as a currently active portfolio. If the currently active portfolio population does not result in an accepted behaviour, such indirect activations may occur to supplement the available information.

Such indirect activations must be managed to ensure their relevance. Indirect activations are therefore behaviours which must have adequate recommendation strengths in the currently active population to be accepted.

When two portfolios are active at the same time, they acquire recommendation strengths in favour of activating each other. This recommendation strength decays with time, but if such activation actually occurs and is followed by positive consequences the decay is reversed. Similarly, if two portfolios record conditions at the same time, they acquire recommendation strengths in favour of activating each other, with the same decay properties.

Such indirect activation provides a straightforward model for word interpretation. For example, the set of portfolios often detected within visual experiences of different instances of an object category like CAT are often active at the same time as the set of portfolios often detected within auditory experiences of the word "cat". The auditory portfolios therefore acquire recommendation strengths in favour of activating the corresponding visual portfolios. These visual portfolios will result in the experience of a visual image. The activated portfolios will be those most often active when the word has been present at the same time as a visual experience, and will therefore be an average of such past visual experiences. The portfolios in arrays closest to sensory inputs will occur less consistently in response to different instances of any one category, and will therefore not be activated. The pseudovisual experience will therefore not be an hallucination.

A portfolio can acquire recommendation strength in favour of identifying a particular category. Such acquisition may be random or aided by genetically defined imitation behaviours [3]. The appropriate visual portfolios will therefore have recommendation strengths in favour of saying the word "cat". Note that any one portfolio may have recommendation strength in favour of identifying a number of different categories, high integrity identifications are only achieved across the currently active portfolio population.

5. PHYSIOLOGICAL RESEMBLANCES TO THE RECOMMENDATION ARCHITECTURE

As described in detail elsewhere [3; 5] there are a number of resemblances between the forms required by the recommendation architecture and the physiological structure of the brain. The separation in the mammal brain between cortex and various subcortical structures including thalamus, basal ganglia and cerebellum resembles the recommendation architecture separation between clustering and various competitive subsystems in a number of ways..

Clustering must be organized into layers with columns extending across several layers and areas made up of parallel columns, strongly resembling cortex organization. Devices in clustering must be organized to detect activity in different groups of inputs corresponding with different conditions. The organization of pyramidal cortex neurons into different dendrites, each with a group of inputs, resembles the required form of clustering devices.

Portfolios as discussed in this paper would therefore be instantiated in cortex columns.

6. MODELLING CATEGORIZATION IN THE RECOMMENDATION ARCHITECTURE

Six different types of phenomena related to categorization were discussed in section 3. The first was categorization on the basis of similarity. The second was typicality effects such as faster categorization of more typical instances. The third was the influence of recently perceived category instances on categorization accuracy. The fourth was rule based categories, the fifth ad hoc categories and the sixth theory based categories. The way in which direct and indirect activation of portfolios supports all these phenomena will now be described.

In the case of similarity based categorization, visual portfolios will be activated in response to the category name and in response to the instance. Categorization depends upon detection of the degree of overlap between the two populations of activated portfolios. A typical category member will result in large overlap, a non member small overlap, and such categorization decisions will be rapid.

The presence of an activated portfolio in two separate populations is itself a condition which could be instantiated in a higher level portfolio. Such portfolios would detect overlap and could acquire recommendation weight in favour of saying "yes, the instance belongs to the category". Such recommendation weights could be modified by consequence feedback where the consequence feedback derived from a supervisor.

For moderate overlap a second phase of activation is required. If the instance is a member of the category but atypical, the overlap will be moderate, but overlap may also be moderate for a non member with some similarity to the category. However, in the case of an atypical category member, many of the portfolios ac-

tivated in response to the atypical member will also have been activated in response to other instances of the category in the past. Hence if the population of portfolios activated in response to the atypical instance is supplemented by portfolios often active in the past at the same time, the overlap with the category name portfolio population will be significantly higher. Indirect activation on the basis of frequent simultaneous past activity can therefore provide discrimination between atypical category members and accidentally similar non members. The time required for the indirect activation accounts for the observed typicality effects.

The influence of recent instances on categorization accuracy follows from the fact that the ability of two portfolios to activate each other indirectly on the basis of simultaneous activity decays with time. In the dermatological example, there will be a group of visual portfolios that tend to be activated in response to hearing the name of a particular condition. There will be another group of portfolios which was activated in response to a recent instance of the condition. The two groups were active at the same time and in general recorded conditions at the same time because the subjects were told the name of the condition associated with the instance. Portfolios in one group will therefore have acquired relatively strong recommendation strengths in favour of indirect activation of portfolios in the other group. If a similar instance is presented, the group of portfolios activated will have significant overlap with group activated in response to the earlier instance. Portfolios in the overlap will have enhanced recommendation strength in favour of expanding the group with the recently simultaneously active portfolios in the category group. Hence the overlap between the expanded group and the category group will be increased by the recent similar instance, increasing the probability of an accurate categorization.

Consider now the support for ad hoc categorization. The example described was THINGS TO TAKE ON A CAMPING TRIP. The mechanism here is that the word "camping" and the word for a possible thing both activate populations of pseudovisual portfolios. If many of the portfolios in one population recorded conditions in the past at the same time as portfolios in the other population, the categorization is confirmed. Alternatively, if a subject is asked to suggest instances of the category, portfolios activated by the word "camping" can indirectly activate other portfolios which recorded conditions in the past at the same time. These indirectly activated portfolios will contain portfolios directly activated in response to articles used while camping, and such portfolios will have recommendation strengths in favour of speaking the names of the articles.

In the example of a dolphin with the appearance of a fish but the internal structure of a mammal, the word "dolphin" will have been present in the past at similar times to portfolios activated in response to illustrations and discussion of the internal structure of

dolphins. Many of these portfolios will also have been activated at the same time as the word “mammal”. The portfolio populations activated by the words “fish” and “mammal” will contain portfolios corresponding with both visual appearance and internal structure. The word “dolphin” will therefore generate a population which has significant overlap with both FISH and MAMMAL portfolio populations. Category assignment will therefore depend upon the relative weight given to overlaps of different types.

In the case of the triangle with one corner cut off, again there are overlaps between the portfolio population activated in response to perception of the figure and both the population activated in response to the word “triangle” and the population activated in response to the word “quadrilateral”. The relative weight given to the different overlaps will determine the categorization selected.

In the recommendation architecture model, category membership is thus defined by the overlap between the portfolio population activated in response to the category name and the population activated in response to the instance, with the instance population supplemented by portfolios often active in the past at the same time as the directly activated instance portfolios, and the category name population sometimes supplemented by portfolios which recorded conditions in the past at the same time as the directly activated category portfolios.

7. CONCLUSIONS

The recommendation architecture is the form into which any system which learns a complex combination of behaviours will be constrained if resources are not unlimited. There are strong resemblances between the physiology of the brain and the recommendation architecture forms. The information storage and access mechanisms in a system with the recommendation architecture provide a natural account for the complete range of observed human categorization phenomena. Only one form of information storage is needed to support all the categorization phenomena. The different phenomena follow from different ways in which the same information is accessed in the recommendation architecture.

8. REFERENCES

- [1] Coward, L. A. (2004). The Recommendation Architecture Model for Human Cognition. Proceedings of the Conference on Brain Inspired Cognitive Systems, University of Stirling, Scotland.
- [2] Gedeon, T., Coward, L. A., and Zhang, B. (1999). Results of Simulations of a System with the Recommendation Architecture, Proceedings of the 6th International Conference on Neural Information Processing, Volume I, 78-84.
- [3] Coward, L. A. (2001). The Recommendation Architecture: lessons from the design of large scale electronic systems for cognitive science. *Journal of Cognitive Systems Research* 2(2), 111-156.
- [4] Ratnayake, U. and Gedeon, T. D. (2002). Application of the recommendation architecture for discovering associative similarities in documents. *Proceedings of the 9th International Conference on Neural Information Processing*.
- [5] Coward, L. A. (2005). *A System Architecture Approach to the Brain: from Neurons to Consciousness*. New York: Nova Science Publishers.
- [6] Wittgenstein, L. (1958). *Philosophical Investigations*. Oxford: Blackwell.
- [7] Rips, L., Shoben, J. and Smith, E. (1973). Semantic Distance and the Verification of Semantic Relations. *Journal of Verbal Learning and Verbal Behaviour* 12, 1 – 20.
- [8] Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General* 104, 192 – 232.
- [9] Hampton, J. A. (1995). Testing the prototype theory of concepts. *Journal of Memory and Language* 34, 686 – 708.
- [10] McCloskey, M. and Glucksberg, S. (1978). Natural categories: well-defined or fuzzy sets? *Memory and Cognition* 6, 462 – 472.
- [11] Nosofsky, R. M., & Johansen, M. K. (2000). Exemplar-based accounts of “multiple-system” phenomena in perceptual categorization. *Psychonomic Bulletin & Review* 7, 375 - 402.
- [12] Nosofsky, R. M., Kruschke, J. K., & McKinley, S. C. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 211–233.
- [13] Brooks, L. R., Norman, G. R. and Allen, S. W. (1991). Role of specific similarity in a medical task. *Journal of Experimental Psychology: General* 120 (3), 278 – 287.
- [14] Minda, J. P. and Smith, J. D. (2002). Comparing prototype-based and exemplar-based accounts of category learning and attentional allocation. *Journal of Experimental Psychology: Learning, Memory and Cognition* 28 (2), 275 – 292.
- [15] Barsalou, L. W. (1983). Ad hoc categories. *Memory and Cognition*, 11, 211-227.
- [16] Tiensten, J. L. (1988). Resemblance and general terms. *Philosophical Studies* 54, 87 – 108.
- [17] Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge MA: MIT Press.
- [18] Sloman, S. A. and Rips, L. J. (1998). Similarity as an explanatory concept. *Cognition* 65, 87 – 101.
- [19] Keil, F. C., Smith, W. C., Simons, D. J. and Lev in, D. T. (1998). Two dogmas of conceptual empiricism: Implications for hybrid models of the structure of knowledge. *Cognition* 65, 17-49.
- [20] Smith, E. S., Patalano, A. L. and Jonides, J. (1998). Alternative strategies of categorization. *Cognition* 65, 167 – 196.
- [21] Hyvärinen, A., Karhunen, J. and Oja, E. (1999). *Independent Component Analysis*. New York: Wiley.