Adaptive and Holistic Knowledge Representations Using Self-Organizing Maps

Timo Honkela
University of Art and Design, Media Lab
Hämeentie 135 C, FIN-00560 Helsinki, Finland
e-mail: timo.honkela@uiah.fi, fax: + 358 9 7563 0555

Abstract
This paper discusses the need for adaptivity and gradience in knowledge representation systems. Problems related to predetermined classification systems and to use of fixed semantic primitives are presented. The qualities of the Kohonen's self-organizing map algorithm are considered as knowledge representation formalism useful especially in complex and conceptually dynamic domains.

Keywords
Self-Organizing Map, SOM, knowledge representation, unsupervised learning

Introduction

When written language is analyzed, the static and discrete features of natural language are easily overemphasized. For instance, the majority of knowledge representation formalisms used in the artificial intelligence and natural language processing research have been based on the tacit assumption of the world consisting of entities and relations between these entities. The words in the language are considered to be "labels" of the entities. The formalisms based on predicate logic contain predicates and their arguments (to model static structures of entities and their relations), logical connectives and quantifiers, and implications (to model rule-like phenomena and dependencies). Semantic networks and frame systems share the underlying ontological assumption. The influence of the "classical AI techniques" on the models and metaphors used in cognitive science is substantial.

Handling gradience

Natural language is nowadays usually not viewed as a means for labeling the world but being an instrument by which the society and the individuals within it construct a model of the world. The world is continuous, and changing. Thus, the language is a medium of abstraction rather than a tool to create an exact "picture" of selected portions of the world. In the abstraction process, the relationship between a language and the world is one-to-many in the sense that a single word or expression in language is most often used to refer to a set or to a continuum of situations in the world. Thus, in order to model the relationship between language and world, the apparatus of the predicate logic, for instance, does not seem to provide enough representational power.

One way of enhancing the representation is to take into account the unclear boundaries between different concepts. Many names have been used to refer to this phenomenon such as gradience, fuzziness, impreciseness, vagueness, or fluidity of concepts.

Need for adaptation

However, handling gradience not enough because the system has to be able to adapt to the different conceptual systems of the other, and the development and changes in the domains being discussed. In other words, concepts may be considered as areas in the domain space with fuzzy boundaries but, moreover, the areas and the links between the concepts and the terms that refer to them are individually determined by each person in a long learning process. Thus, taking this kind of epistemologically relativist position, one may state that no two persons have exactly the same conception of (meaning for) of any word in a language. The consequences of such an idea may seem to be radical at the first glance: how is communication possible in the first hand? Does this also mean that anything goes and there are no limits? No, that is not the case, either.

When one considers epistemological questions (questions of knowing) in the framework of classical logic problems in such a view seem to be unsolvable: either you believe in something or you believe in its opposite. However, when one adopts the background assumptions of the classical logic one also tends to consider the world as a collection of entities and their relationships. Ontologically propositions like "snow of white" or "the pine grows in my yard" refer to highly complex states of the world with possibly millions of details (some of which are relevant in evaluating the propositions and some not). Moreover, there are a large number of more or less different instances of "snow", "white things", "pines", "yards", etc. There is a one-to-many mapping the words to the world. Ambiguities make this mapping even more complicated. Because one "state of the world" can be described in several ways the mapping becomes many-to-many.

The contents of propositions such as shown above should not be conceptualized straightforwardly as, e.g., in the model-theoretical approach (and related approaches) in logic or in the knowledge representation formalisms of artificial intelligence or cognitive science. These simplifications may be often useful but problems arise if the simplistic conceptualization is applied without considering the complex relationship between the perceptions and the languages, and the conceptualization of the whole domain formed in the speaker community

In order to develop systems that would be able to communicate with human beings in fine-grained and contextually tuned manner, one has to develop the ways how machines can experience the world (e.g., through images), learn from the experience, and learn to associate language with the perceptions, or more accurately, with the perceptional clusters, and finally create conceptual systems with a reference to the languages and conceptual systems of human beings (ranging from languages to individuals).

Self-Organizing Maps

Kohonen's Self-Organizing Map (SOM) [1] is a means for automatically arranging high-dimensional statistical data. The map attempts to represent the input samples with optimal accuracy using a restricted set of models. The models also become ordered on the map grid in such a way that similar models appear close to each other and dissimilar models far from each other. The SOM is useful in clustering, abstraction, and visualization through dimensionality reduction. The unsupervised learning scheme of the SOM makes it suited for applications in which the input data cannot be labeled.

The majority of the thousands of applications in which the SOM has been used have been based on analysis and visualization of numeric data. However, also symbolic and linguistic data has been analyzed. Ritter and Kohonen [2] created the first semantic maps in which the words were organized automatically into the map based on the context of the word. Those words that appeared in similar contexts in the automatically generated input text are positioned close to each other on the map. In a more recent study [3], a real text corpus of Grimm fairy tales was analyzed. On the resulting map the verbs and nouns formed areas of their own on the opposite sides of the map. Moreover, the area of the nouns became divided into subareas of animate and inanimate nouns.

Using the context in the text is one potential source for creating input for a map of words [4,5]. Also other kinds of contexts are possible. For instance, an adjective map shown in Figure 1 is based on the evaluation of each word on a 12-dimensional scale. Each dimension referred to an emotive aspect that was evaluated. However, the map below serves only as an illustration as it is based on the judgements of only one person. Studies can be made in which the subjectivity and intersubjectivity is analyzed and modeled on a map.

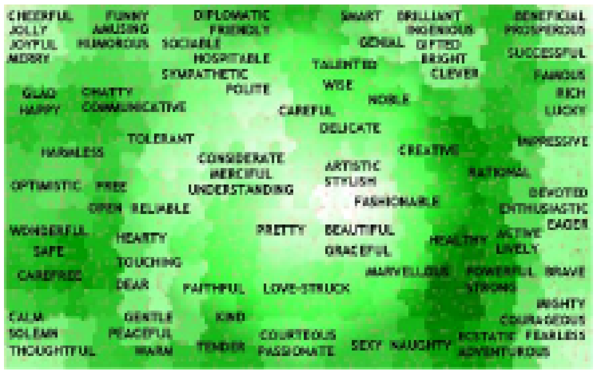

Figure 1. A collection of adjectives organized by self-organizing map. Each adjective has been associated with 12 parameters that characterize the emotive nature of the word.

Problems of categorization

One can consider problems related to knowledge representations useful in descriptions of audiovisual data. If pictorial or sound data is considered a classification or textual description is usually needed for finding pieces of data. Often the description is based on a pre-defined classification or a list of keywords, i.e. a terminology base or on a thesaurus. However, even if the identity of the artist or the place of publishing can be rather easily determined unambiguously, the same is not true for the description of the contents. For instance, in the domain of information retrieval and databases of text documents, Furnas et al. [6] have found that in spontaneous word choice for objects in five domains, two people favored the same term with less than 20% probability. Bates [7] has shown that different indexers, well trained in an indexing scheme, might assign index terms for a given document differently. It has also been observed that an indexer might use different terms for the same document at different times. Moreover, the conceptually constructed reality in any active domain changes over time and thus classification systems created earlier may become obsolete. This phenomenon is visualized in Fig 2. The example can be considered to be relevant when the attempts to create fixed semantically oriented representational systems are considered.

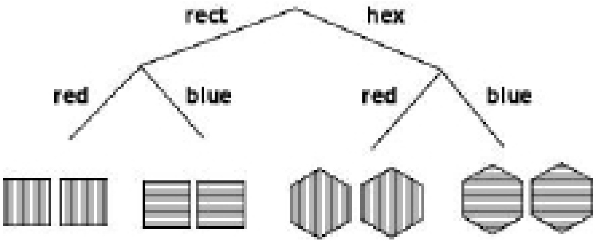

Figure 1. Illustration of a set of unproblematic cases in relation to a fixed classification systems. However, over the time some of the categories may not be needed anymore, and some new ones would be needed.

Even when no classification is used but the information retrieval is based on use of key words and full-text documents problems arise. The traditional key word based approach with Boolean logic has three basic problems. First, for Boolean queries there is no simple way of controlling the size of the output, and the output is not ranked in the order of relevancy. In addition, considering the results of a query it is not known what was not found, especially if the collection is unfamiliar. Third, if the domain of the query is not known well it is difficult to select the appropriate key words. Thus, even if the indexer or the metadata creator is able to find accurate descriptions of the content, the user of the metadata may not succeed in that, i.e. to use the same words or phrases.

It is a very basic problem in text document management that different words and phrases are used for expressing similar objects of interest. Natural languages are used for the communication between human beings, i.e., individuals with varying background, knowledge, and ways to express themselves. When audiovisual contents are considered this phenomenon should be more than evident.

Adaptive prototypes

In the following, an illustrative example is given to show the self-organizing map can be used to handle a situation in which the conceptual basis of the data is changing. The basic idea is that the self-organizing map modifies itself to model the input as faithfully as possible. The map elements or vectors can be considered to be adaptive prototypes that represent the domain in a dynamic manner. This can be contrasted with the approach in which the classification system is fixed or a fixed set of semantic primitives is used to model the data.

Figure 2 illustrates the analysis of a set of items in which each item belongs strictly to one of the three color categories and one of the three shape categories.

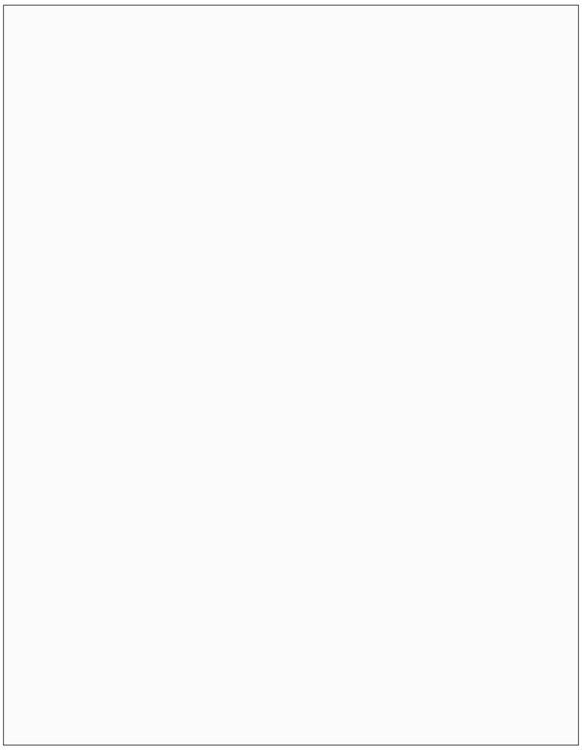

Figure 2. Map of 30 items. Each of them is strictly red, blue or green in color, and triangular, squared or hexagonal in shape. The labels indicate the index of each item, the color and the shape.

The "coordinates" over the map seem to become formed based on the most common cases of red triangles, green triangles and green squares. However, this kind of analysis might be of little use if the data items would remain such that they could be easily positioned into the given classes. Figure 3 illustrates an analysis of a data set that had 10 items in common with the earlier one. The rest 20 items consisted of objects that could have any color (in 3-dimensional real-valued vector of RGB values) and the number of edges varied from 3 to 9. The map defines 20 prototypes of items that are organized in such a way that those two items that resemble each other in the input space tend to appear close to each other on the map. In the latter example (shown in Figures 3 and 4) this neighborhood relation is more useful because there are mediating cases between the most prototypical ones. Moreover, the map also visualizes the existence and proportion of the “non-prototypical” cases, e.g. the items here that do not belong to the categories of red, blue or green.

Figure 3. Map of 30 items: 10 items drawn from the set visualized in Figure 2, and 20 items as examples of conceptual change and intercategory cases.

Conclusion

One may ask why is it relevant to present simple illustrations of the use of the self-organizing map such as the ones of this article. However, it seems that the basic ideas presented in this article and earlier in many other contexts (see, e.g., [2]) have not been taken into account when knowledge representation formalisms are being developed. This appears to be true, for instance, in the areas of ontological engineering and the use of XML. The real-world application needs related to the representation of changing knowledge based on intersubjective “agreements” seems to point out the need of using adaptive prototypes rather than fixed categories or semantic primitives.

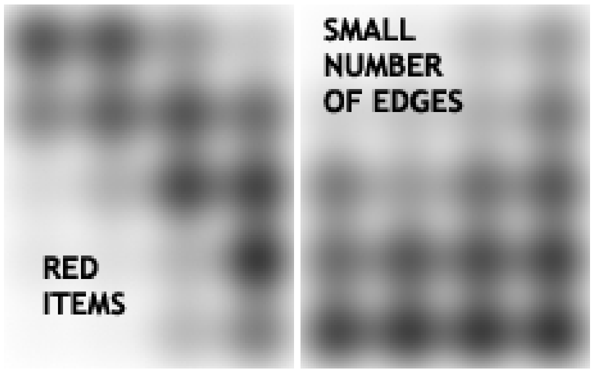

Figure 4. Component levels showing the distribution of the redness (on the left) and the feature denoting the number of edges of the items (on the right). The order of the map is the same as in Figure 3

References

1. Kohonen, T. (1995). Self-Organizing Maps. Berlin & Heidelberg: Springer.
2. Ritter, H. & Kohonen, T. (1989). Self-organizing semantic maps. *Biological Cybernetics*, 61, 4, 241-254.
3. Honkela, T., Pulkki, V., & Kohonen, T. (1995). Contextual Relations of Words in Grimm Tales Analyzed by Self-Organizing Map. *Proceedings of ICANN-95, International Conference on Artificial Neural Networks*, Volume 2, F. Fogelman-Soulié and P. Gallinari (eds.), EC2 et Cie, Paris, pp. 3-7.
4. Honkela, T. (1997). Learning to understand: general aspects of using self-organizing maps in natural language processing. In D. Dubois (ed.), *Proceedings of the CASYS'97, Computing Anticipatory Systems*, Woodbury, New York: American Institute of Physics, pp. 563-576.
5. Honkela, T. (2000). Self-organizing maps in symbol processing. *Hybrid Neural Systems*, Stefan Wermter, Ron Sun (eds.), Springer, Heidelberg.
6. Furnas, G. W., Landauer, T. K., Gomez, L. M., & Dumais, S. T. (1987). The vocabulary problem in human-system communication. Communications of the ACM, 30, 11, 964-971.
7. Bates, M. J. (1986). Subject access in online catalog: a design model. Journal of the American Society of Information Science, 37, 6, 357-376.