

# Von Foerster meets Kohonen - Approaches to Artificial Intelligence, Cognitive Science and Information Systems Development

**Timo Honkela**

Helsinki University of Technology  
Neural Networks Research Center  
Espoo, Finland  
timo.honkela@hut.fi

## Abstract

In this article, some aspects of Heinz Von Foerster's work that are of particular importance for cognitive science and artificial intelligence are studied. Kohonen's self-organizing map is presented as one method that may be useful in implementing some of Von Foerster's ideas. The main foci are the distinction between trivial and non-trivial machines and the concept of constructive learning. The self-organizing map is also presented as a potential tool for alleviating the participatory crisis discussed by Von Foerster. In the end, the importance of considering change in information systems development is considered.

**Keywords:** artificial intelligence, cognitive science, natural language processing, self-organizing map, machine learning, emergence.

# 1 Introduction

In this article, I will outline some aspects of Heinz Von Foerster's work that are of particular importance for cognitive science, artificial intelligence and information systems development. The general consideration is preceded by an introductory example based on experiences as an individual researcher. Namely, I began my work as a research assistant in 1984 in a large project that was developing a natural language database interface for Finnish. Methodologically the project (Kielikone, "Language Machine", funded by Sitra foundation) was following the traditional artificial intelligence approach that was popular in 1980's (consider, e.g., Jäppinen and Ylilammi, 1986, Jäppinen et al., 1988).

Systems for natural language processing as well as expert systems were commonly based on a collection of rules or some other symbolic representations. The architecture of this particular system was ambitiously designed to cover the levels of written language: lexicon, morphology, syntax, semantics, and even pragmatics. Modeling of the structural level of Finnish language, i.e, morphology and syntax, was successful (Jäppinen and Ylilammi, 1986, Valkonen et al., 1987). Finnish has complex morphology which can be exemplified by the fact that each verb may have about 12,000 inflectional forms (Karls-son, 1999). The rule-based system for morphological analysis reached a level comparable to the skill of a native speaker of Finnish (Jäppinen and Ylilammi, 1986). Another high-quality system for morphological analysis was developed by Koskenniemi (1983) with a general two-level representation system applicable for a large number of languages.

In the Kielikone project, the syntactic analyzer was based on a dependency grammar: the words in a sentence are taken to depend on each other and, thus, create a hierarchical structure. The dependency grammar fits well with the quality of Finnish language where the word order is relatively free. An alternative approach is, e.g., to develop phrase structure grammars with categories that correspond to a whole sentence or some parts of it such as verb phrases, nominal phrases, nouns, and verbs. The rules of a phrase structure grammar are based on the idea of fixed word order in which, for instance, the subject precedes the main verb of the sentence like in English.

A natural language database interface is supposed to analyze the meaning,

not only the structure of the input sentences. When associated with a stock exchange data base the system should be able to respond properly to expressions such as “Show me the largest companies in forestry!”, or “How many companies have turnover larger than 1 billion euros?”. Some users would also ask questions that require analysis of imprecise or fuzzy expressions like “What are the companies with a large turnover and a small number of employees”, or even prediction: “Which companies will be profitable next year?”. The user expectations are relatively high if the system is coined with the term intelligent. In our project, the semantic analysis was conducted using a rule-based system that was transferring the results of the syntactic analysis, i.e., dependency trees into predication structures which were variants of expressions in predicate logic. All the basic assumptions of predicate logic as a representation formalism were present. One of the central ideas is that the world, or domain of interest, can be adequately modeled as a collection of objects and relationships between them. This basic idea and the notion of hierarchical conceptual relations is also prominent in semantic nets as well as in object-oriented design and programming in computer science.

The efforts in developing the natural language interface in the Kielikone project as well as in many others before and after this one have shown that the acquisition of knowledge sufficient for in-depth understanding of a variety of questions and commands is highly difficult. Considerable success has been gained only in systems that have a limited domain of application. A classic example is the SHRDLU system by Winograd (1972). It was able to process commands related to a collection of items on the screen with varying sizes, colors and shapes. It is also important that the user is well aware of the domain and intended use of the system. Similarly, in machine translation systems (Hutchins, 1986, Nirenburg et al., 1992) it has been commonplace to limit the application to a specific area such as weather forecasts (Chandioux, 1976, Isabelle, 1987) or product descriptions (Lehtola et al., 1999).

In the following, I will outline how becoming familiar with Heinz von Foerster’s research made it possible to consider critically the approach used earlier and to apply and develop alternative methods and approaches. In the end, I will discuss some topics with the aim to show that, even though Von Foerster’s thinking is becoming widely recognized among researchers, there are still many aspects of his results that deserve wider attention and have further potential in practical applications.

## 2 Observing Systems

Let's first consider one quote from "Notes on an epistemology for living things" that was originally published in 1972: "Objects and events are not primitive experiences. Objects and events are representations of relations. Since 'objects' and 'events' are not primary experiences and thus cannot claim to have absolute (objective) status, their interrelations, the 'environment' is a purely personal affair, whose constraints are anatomical or cultural factors. Moreover, the postulate of an 'external (objective) reality' disappears to give way to reality that is determined by modes of internal computations." (Von Foerster, 1981a, p. 261) When computational models of the semantics of natural language are built using predicate logic the status of the predicates is often taken as granted. This corresponds to the idea of defining that the sentence "snow is white" is true if and only if snow is white. This kind of consideration of pure logical form and assuming a rather straightforward one-to-one correspondence between language and the world are problematic which is shown clearly by Von Foerster. For instance, he mentions that "a formalism necessary and sufficient for a theory of communication must not contain primary symbols representing communicabilia (e.g. symbols, words, messages, etc.)" (Von Foerster, 1981a, p. 262/267). Gärdenfors (2000) explains in a detailed manner why it is important to consider the level of observation when building a theory of semantics and conceptual systems. Also Zadeh has emphasized the importance of the observation (Zadeh, 2002). Formal logicians wish to consider semantics without reference to any cognitive agent or observer, By contrast, Von Foerster (1981a, p. 263) points out that "'information' is a relative concept that assumes meaning only when related to the cognitive structure of the observer of an utterance (the 'recipient')."

For many ideas presented by Von Foerster, I found an operational context when I familiarized myself with Teuvo Kohonen's work, in particular his self-organizing map in the end of 1980s and in the beginning of 1990s. However, I am not aware of any explicit link between Von Foerster's and Kohonen's work.

In the following, learning systems in general, and specifically the self-organizing map as well as its application potential in the areas of cognitive science and artificial intelligence are described in some detail.

## 2.1 Learning systems

In the area of machine learning in artificial intelligence, the principles of adaptation in natural systems have been studied and similar principles have been implemented as computational systems. A fundamental reason for adaptation is described by Von Foerster (1981b, p. 196): “At any moment we are free to act toward the future we desire. In other words, the future will be as we wish and perceive it to be. This may come as a shock only to those who let their thinking be governed by the principle that demands that only the rules observed in the past shall apply to the future. For those the concept of ‘change’ is inconceivable, for change is the process that obliterates the rules of the past.”

If a system is supposed to function properly in changing conditions, it either needs to be constantly reprogrammed or it must be adaptive. Machine learning can be divided into three categories: it can be supervised, reinforced and unsupervised.

## 2.2 Supervised and unsupervised learning

In supervised learning, the system is given input-output pairs: for each input there must also exist the “right answer” to be enforced at the output. The system then learns these input-output pairs. The task is not trivial, however, and after the learning period the network is also able to deal with inputs that were not present in the learning phase. This property ensues from the generalizing capabilities of the network. Supervised learning is often used for classification tasks.

Among neural network models, the most widespread supervised learning method is the backpropagation algorithm (cf., e.g., Rumelhart et al., 1986). The drawback of supervised learning is the need for correct output. In some cases, obtaining the output for each input case is a very laborious task if large source material is used. It may sometimes even be impossible to determine a unique or correct output, by definition, for a certain input.

Moreover, the output is determined using some ready-made classification system. von Foerster points out that “we seem to be brought up in a world seen through descriptions by others rather than through our own perceptions.

This has the consequence that instead of using language as a tool with which to express thoughts and experience, we accept language as a tool that determines our thoughts and experience.” (Von Foerster, 1981b, p. 192) One can characterize learning (“bringing up”) in a supervised manner as being based on descriptions by others.

In reinforced learning, the system is given an estimate of how good the result generated was, rather than the correct output.

Whereas supervised learning models are suitable for classification, unsupervised learning can be used for abstraction. The self-organizing map, applying an unsupervised learning principle, enables autonomous processing. In the following, the self-organizing map is considered in detail.

### 3 Self-Organizing Map

The self-organizing map (Kohonen 1982, 1993, 2001) is a widely-used artificial neural network model in which learning is unsupervised: no a priori classifications for the input examples are needed. Related research have been conducted, for instance, by Amari (1980), Carpenter and Grossberg (1991), and von der Malsburg (1973). The network architecture of the self-organizing map consists of a set of laterally interacting adaptive processing elements, nodes, usually arranged as a two-dimensional grid called the map. All the map nodes are connected to a common set of inputs. Any activity pattern on the input gives rise to excitation of some local group of map nodes. After learning, the spatial positions of the excited groups specify a mapping of the input onto the map. The learning process is based on similarity comparisons in a continuous space. The result is a system that map similar inputs close to each other in the resulting map. The input may be highly complex multidimensional data like in real-life applications, such as, speech recognition (e.g., Kohonen, 1988), image analysis (e.g., Visa and Iivarinen, 1997), recognition of handwritten characters (e.g., Vuori et al., 2001), decision support (e.g., Carlson, 1991), financial analysis (e.g., Back et al., 1996, DeBoeck and Kohonen, 1998), information retrieval (e.g., Lin et al., 1991, Merkl, 1994, Honkela et al., 1996, Kaski et al., 1998, Lagus et al., 1999) and process monitoring (e.g., Kohonen et al. 1996, Simula et al. 1999). The theoretical aspects of the algorithm have also been studied extensively (see, e.g., Erwin

et al. 1992). An extended version, adaptive subspace self-organizing map, is able to deal with variation in the input and find invariant features (Kohonen, 1996, Kohonen et al., 1997).

### 3.1 Learning algorithm

Starting with an initially random set of prototypes, the self-organizing map algorithm gradually adjusts them to reflect the clustering of the training data. In the following, this process is explained in detail.

Assume that some sample data sets have to be mapped onto the array depicted in Fig. 1. A sample set is described by a real vector  $x(t) \in R^n$  where  $t$  is the index of the sample, or the discrete-time coordinate. In setting up the self-organizing map, one first assigns to each unit in the array a parameter vector  $m_i(t) \in R^n$  called the prototype vector, which has the same number of elements as the input vector  $x(t)$ . The initial values of the parameters (components of  $m_i(t)$ ) can be selected at random. The process described below changes these parameters. (Kohonen 2001)

The “image” of an input item on the map is defined to be in the location, the  $m_i(t)$  which matches best with  $x(t)$  in some metric. The self-organizing algorithm that creates the ordered mapping can be described as a repetition of the following basic tasks:

- An input vector  $x(t)$  is compared with all the prototype vectors  $m_i(t)$ . The best-matching unit on the map, i.e., the unit where the parameter vector is most similar to the input vector in some metric, called the winner, is identified.
- The prototype vectors of the winner and a number of its neighboring units in the array are changed incrementally according to the learning principle specified below.

The basic idea in the self-organizing map is that for each input sample vector  $x(t)$  the parameters of the winner and units in its neighborhood are changed closer to  $x(t)$ . For different  $x(t)$  these changes may be contradictory, but the net outcome in the process is that ordered values for the  $m_i(t)$  are finally

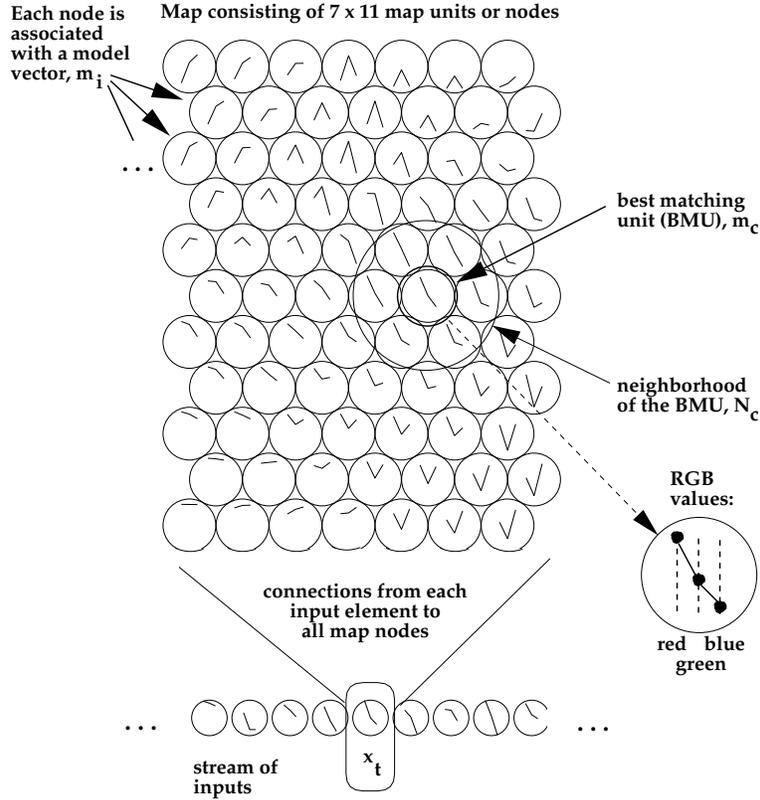


Figure 1: The basic architecture of Kohonen's self-organizing map.

obtained over the array. If the number of input vectors is not large compared with the number of prototype vectors (map units), the set of input vectors must be presented reiteratively many times. As mentioned above, the prototype vectors may initially have random values, but they can also be selected in an ordered way. Adaptation of the prototype vectors in the learning process takes place according to the following equation:

$$m_i(t+1) = m_i(t) + \alpha(t)[x(t) - m_i(t)] \quad \text{for each } i \in N_c(t), \quad (1)$$

where  $t$  is the discrete-time index of the variables, the factor  $\alpha(t) \in [0, 1]$  is a scalar that defines the relative size of the learning step, and  $N_c(t)$  specifies the neighborhood around the winner in the map array. At the beginning of the

learning process the radius of the neighborhood is fairly large, but it shrinks during learning. This ensures that the global order is obtained already at the beginning, whereas towards the end, as the radius gets smaller, the local corrections of the prototype vectors in the map will be more specific. The factor  $\alpha(t)$  decreases during learning. (Kohonen 2001)

### 3.2 Self-organizing maps as non-trivial machines

Occasionally connectionist models are claimed to be neobehavioristic. In Von Foerster's text behavioristic systems are called trivial machines: "A trivial machine is characterized by a one-to-one correspondence between its 'input' (stimulus, cause) and its 'output' (response, effect). This invariable relationship is 'the machine'. Since this relationship is determined once and for all, this is a deterministic system; and since an output once observed for a given input will be the same for the same input given later, this is also a predictable system. Non-trivial machines, however, are quite different creatures. Their input-output relationship is not invariant, but is determined by the machine's reactions. While these machines are again deterministic systems, for all practical reasons they are unpredictable: an output once observed for a given input will most likely be not the same for the same input given later. In order to grasp the profound difference between these two kinds of machines it may be helpful to envision 'internal states' in these machines. While in the trivial machine only one internal state participates always in its internal operation, in the non-trivial machine it is the shift from one internal state to another that makes it so elusive." (Von Foerster, 1981b, p. 198) Later he writes: "While our pre-occupation with the trivialization of our environment may be in one domain useful and constructive, in another domain it is useless and destructive. Trivialization is a dangerous panacea when man applies it to himself." (Von Foerster, 1981b, p. 199).

The self-organizing map can be considered as a good example of a "non-trivial machine". It is not a behavioristic model. On the contrary, the internal state of the system influences the behavior and it is changing during the "life" of the map. There are applications, though, in which the map is fixed after an initial learning phase but this is a property of the algorithm's application and not a property of the algorithm itself. One of the reasons why the self-organizing map and other artificial neural network models may appear

behavioristic is that their internal state is difficult to grasp and analyze. For instance, multilayer perceptrons typically are used as black-box systems: only the input-output behavior is considered in applications. However, as cognitive models these systems should be considered only in the continuous learning mode: the relationship between the input and the output is constantly subject to change and, thus, the system is non-trivial. As mentioned earlier, the self-organizing map is even less trivial since there is no classification framework predetermined by the designer of the system.

### 3.3 Modeling Constructive Learning

Constructive learning involves qualitative restructuring and modification of internal knowledge representations, rather than just accumulation of new information in memory. Epistemological theories of knowledge have traditionally been based on predicate logic and related methodologies and frameworks. The basic assumption is that the world consists of objects, events and relationships. The language and the conceptual structures are then supposed to reflect rather straightforwardly this ontological structure. Learning has been seen as a means to memorize the mapping from the epistemological domain (to put it simply: words) onto the ontological domain (objects, events and relationships). This view has been dominant at least partly because of the consistent formalization of the theory through the use of symbolic logic. Moreover, the use of the von Neumann computer as the model or metaphor of human learning and memory has had similar effects and has strengthened the idea of memory as a storage consisting of separate compartments which are accessed separately and which are used in storing and retrieving information more or less as such. (Honkela et al., 2000)

The self-organizing map is dynamic, associative and consists of elements that can also be called adaptive prototypes (Honkela, 2000a). Inputs are not stored as such but comparison is made between the input and the collection of prototypes. The closest prototype of the input is adapted towards the input. The same operation is also conducted for the neighboring prototypes, which gives rise to the topographical order on the map. Thus, the adaptation process in the self-organizing map algorithm is based on the principle that what already exists in the system also influences the learning result. (Kohonen, 2001)

Considered superficially, one could claim that modeling learning phenomena through the use of the self-organizing map would be mentalistic. However, it is possible to construct a model in which a number of autonomous map-based agents interact in such a way that they perform a social construction of knowledge and find intersubjective epistemological agreements (Honkela, 1993, Honkela et al., 2003, Honkela and Winter, 2003).

In the model presented in (Honkela and Winter, 2003), simulated autonomous agents associate linguistic expressions and visual perceptions (cf. Fig. 2). The original situation is such that the expressions that the agents use are random: each agent has, in principle, a language of its own. However, in the beginning the information exchange is contextual, i.e., two communicating agents can both perceive the same item. This is the basis for symbol formation, symbol grounding and transformation from subjective individual languages into intersubjective language shared by the community (consider also Steels, 1996). However, even after convergence into a common language, the agents have a certain level of subjectivity, i.e., the reference relation between the language and the world is not identical between any two agents but is generally close enough in order for useful communication to take place. This individuality of interpretation is a natural phenomenon when the multidimensional and continuous nature of both linguistic expressions and perceptions is taken into account. The individuality of interpretation is considered a problem when meaning and understanding are studied in the framework of symbolic representations and within model theoretical approaches.

Earlier, the emergence of linguistic representations based on the self-organizing map has been studied, e.g., to create word category maps in which those words that appear in similar contexts are located close to each other on the map (consider, e.g., Ritter and Kohonen, 1989, Honkela and Vepsäläinen, 1991, Finch and Chater, 1992, Miikkulainen, 1993, Honkela et al., 1995, Honkela, 1997, Honkela, 2000b).

By emphasizing the pattern nature of language and world, we avoid the idea that some relativism would be a problem as it is in the traditional epistemological research in which language is considered as a collection of propositions and the world consists of a collection of distinct objects and their relationships. This line of thought is inherited rather directly from Von Foerster (1981a) who stated that the logical properties of “invariance” and “change” are those of representations. He gave the expression “the distinct

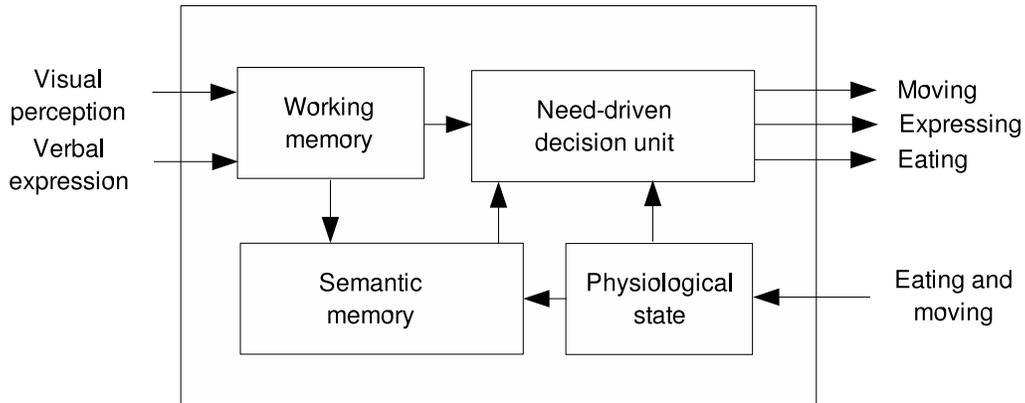


Figure 2: An agent architecture. The agent receives two kinds of perceptual inputs: visual images and linguistic expressions. There are three kinds of potential actions: the agent can either eat, move or utter an expression. The perceptions are primarily stored in the working memory. The semantic memory associates perceptual information and information considering its physiological state. Sudden changes in the physiological state are related to eating actions and the quality of the eaten object determines the direction of the change. The physiological state is also influenced by moving: gradually the agent loses energy. The physiological state serves as the basic motivational factor for the agent. If the energy level is low the agent prefers eating and high energy level makes the exploration of the environment become a more probable action. Communication between the agents is motivated by the exchange of information related to the edibility of the food items. It is assumed that the visual perceptual characteristics of the food items correlate with their level of edibility. The agents do not have explicit information on their position in the environment and therefore their navigation is based on the perceived landmarks in the environment. The environment of an agent consists of other agents, a number of obstacles and food items. Each obstacle has unique visual characteristics which helps the agents in their navigation. See (Honkela and Winter, 2003) for more details.

being the same” as the paradox of invariance, and “the same being distinct” as the paradox of change. Von Foerster expanded the analysis into the nature of objects and events (requote): “Since ‘objects’ and ‘events’ are not primary experiences and thus cannot claim to have absolute (objective) status, their interrelations, the ‘environment,’ is a purely personal affair, whose constraints are anatomical and cultural factors.” (Von Foerster, 1981a, p. 261). In our simulation model (Honkela and Winter, 2003), this principle is exemplified by the fact that both the language and the world model of each agent is purely individual and remains as such. However, there is a certain convergence process that ensures that the agents tend to use the same expressions in similar situations. The sameness in the emerging symbolic level is an invariance of the representation.

## 4 Discussion

In the following, two particular themes related to the intersection of Heinz Von Foerster’s and Teuvo Kohonen’s work are discussed. First, the participatory crisis in society is discussed based on Von Foerster’s article “Responsibilities of Competence”. Second, the concept of change is handled within the framework of information systems development.

### 4.1 Participatory crisis

Von Foerster (1981c, p. 210) characterizes some problematic aspects of modern society in a manner that has become even more relevant after its writing: “It is clear that our entire society suffers from a severe dysfunction. On the level of the individual this is painfully felt by apathy, distrust, violence, disconnectedness, powerlessness, alienation, and so on. I call this ‘participatory crisis’, for it excludes the individual from participating in the social process. The society becomes the ‘system’, the ‘establishment’ or what have you, a depersonalized Kafkanesque ogre of its own will. It is not difficult to see that the essential cause for this dysfunction is the absence of an adequate input for the individual to interact with the society. The so-called ‘communication channels’, the ‘mass media’ are only one-way: they talk, but nobody can talk back. The feedback loop is missing and, hence, the system is out of control.

What cybernetics could supply is, of course, a universally accessible social input device.”

It seems that it would be possible to use the self-organizing map as a component in a social input device. A societal application is to support democratic decision making processes. For instance, one could organize collections of free-form documents into document maps (Honkela et al., 1996, Kaski et al., 1998, Lagus et al., 1999) in order to show the overall opinion landscape on some matter. This approach would ensure, for instance, that predetermined questions do not limit the space of alternatives that could be considered in a decision making process. The self-organizing map also reflects the statistical qualities of the original data: if a phenomenon is common in the data it also occupies a relatively large area on the map. Voting schemes can be seen to have at least two kinds of problems. The questions and provided alternatives determine the space of potential solutions. It is also possible that the form of the question influences greatly the outcome of a vote. To gain insight on some societal issue, one can also analyze some statistical data. While the mapping is nonlinear, exceptional cases or values do not interfere with the analysis but are shown as specific locations in the map.

Following Von Foerster’s ideas, we need ways to enhance communication in society in such a way that bottom-up processes could balance the top-down ones (consider also Von Foerster, 1984). Instead of the television and other mass media, it seems that our society can benefit more from the use of the web and various tools, such as the self-organizing map, that can be used in analyzing the multitude of points of view often hidden in the bottom-up information flows. In the current democratic decision making processes the bandwidth of the bottom-up information flow from the citizens to the decision makers is low. The self-organizing map can be used to analyze and visualize information to facilitate informed decision making. The map can also be used as an “awareness” tool: on a single map one may represent a landscape of a large variety of alternatives that would otherwise remain hidden. Namely, cognitive processes related to decision making are often guided by some predetermined, rather fixed prototypical cases. The self-organizing map and similar methods may be used to find representative prototypes adaptively to facilitate a balanced view on a matter. One need for adaptation stems from continuous change. This aspect is discussed in the following.

## 4.2 Change and Adaptive Information Systems

In the development of information systems, one should remember what Von Foerster has taught about the status and importance of change. Currently, most information systems are developed as “trivial machines” to be predictable and controllable. However, the static nature of the information systems makes them also prone to be “incompatible with the reality”. One reason is that the domain of use is changing. Another, more profound reason is that we all as human beings have individual conceptual systems that we have gained through constructive learning processes. A conceptually static and coarse-grained information system matches with our conceptual systems only partially. This misfit may lead to errors or unjustified procedures. Therefore, it appears necessary that any information system should be adaptive in order to be able to deal with the variety of conceptual construction and in order to be able to conduct meaning negotiations. Another reason for adaptivity is to allow the systems to “forget”: it is not reasonable to keep, for instance, information about somebody’s mistake in a database for an overly extended period of time because the information may lead into a fixation: those who become aware of the information start to consider the person through that information which causes a social trap. Information systems should be able to support “forgetting” and even “forgiving”. Moreover, our information systems should not promote prejudices as rule-based systems, databases, and other similar, conceptually fixed systems easily do: individuals need to be considered as individuals, not as the representatives of the categories into which they are classified.

In summary, many of the ideas introduced by Von Foerster are still relevant and useful for modern researchers and developers within artificial intelligence, cognitive science and information systems development. In this article, I have made an attempt to highlight some of the potential benefits by linking Von Foerster’s work with the research on artificial neural networks and with the self-organizing map, in particular.

## Acknowledgements

I wish to thank reviewers and the editor for their constructive comments on the article. Special thanks are due to Kevin I. Hynnä who provided specific comments and provided editorial assistance. I am also grateful to all the persons who have provided their their time for deepening earlier discussions including but not limited to Dr. Catherine Bounsaythip, Prof. Peter Gärdenfors, Dr. Aapo Hyvärinen, Prof. Heikki Hyötyniemi, Prof. Mauri Kaipainen, Prof. Fred Karlsson, Prof. Samuel Kaski, Prof. Emer., Academician Teuvo Kohonen, Prof. Kimmo Koskenniemi, Ms. Tarja Knuuttila, Dr. Krista Lagus, Prof. Kirsti Lonka, Prof. Dieter Merkl, Prof. Risto Miikkulainen, Prof. Erkki Oja, Mr. Ville Tuulos, Dr. Harri Valpola, Dr. Petri Vasara, Prof. Lotfi Zadeh, and Prof. Jordan Zlatev. I regret the fact that I did not make strong enough effort to find a chance to meet Heinz Von Foerster personally. However, I wish that I have a chance to collaborate people who share Heinz' vision of a better world, however we desire to define the concept of good.

## References

- Amari, S. (1980), "Topographic organization of nerve fields", *Bulletin of Mathematical Biology*, vol. 42, pp. 339-364.
- Back, B., Sere, K., and Vanharanta, H. (1996), "Data mining accounting numbers using self-organizing maps" in Alander, J., Honkela, T., and Jakobsson, M. (eds.), *Proceedings of STeP'96*, Finnish Artificial Intelligence Conference, Finnish Artificial Intelligence Society, Vaasa, Finland, pp. 35-47.
- Carlson, E. (1991), "Self-organizing feature maps for appraisal of land value of shore parcels" in Kohonen, T., Mäkisara, K., Simula, O., and Kangas, J. (eds.), *Proceedings of ICANN'91*, International Conference on Artificial Neural Networks, volume II, North-Holland, pp. 1309-1312.
- Carpenter, G. and Grossberg, S. (1991), *Pattern Recognition by Self-Organizing Neural Networks*, MIT Press, Cambridge, MA.
- Chandioux, J. (1976), "MÉTÉO: un système opérationnel pour la traduction automatique des bulletins méereologiques destinés au grand public", *META*,

vol. 21, pp. 127-133.

Deboeck, G. and Kohonen, T. (eds.) (1998), *Visual Explorations in Finance with Self-Organizing Maps*, Springer.

Erwin, E., Obermayer, K. and Schulten, K. (1992), "Self-organizing maps: ordering, convergence properties and energy functions", *Biological Cybernetics*, vol. 67, pp. 47-55.

Finch, S. and Chater, N. (1992), "Unsupervised methods for finding linguistic categories" in Aleksander, I. and Taylor, J. (eds.), *Artificial Neural Networks*, vol. 2, North-Holland, pp. 1365-1368.

Gärdenfors, P. (2000), *Conceptual Spaces*. MIT Press.

Honkela, T. and Vepsäläinen, A.M. (1991), "Interpreting imprecise expressions: Experiments with Kohonen's self-organizing maps and associative memory" in Kohonen, T., Mäkisara, T., Simula, O. and Kangas, J. (eds.), *Proceedings of ICANN'91, International Conference on Artificial Neural Networks, Volume I*, North-Holland, pp. 897-902.

Honkela, T. (1993), "Neural Nets that Discuss: A General Model of Communication Based on Self-Organizing Maps" in Gielen, S. and Kappen, B. (eds.), *Proceedings of ICANN'93, International Conference on Artificial Neural Networks*, Springer-Verlag, pp. 408-411.

Honkela, T., Pulkki, V., and Kohonen, T. (1995), "Contextual Relations of Words in Grimm Tales Analyzed by Self-Organizing Map" in Fogelman-Soulié, F. and Gallinari, P. (eds.), *Proceedings of ICANN-95, International Conference on Artificial Neural Networks, Volume II*, EC2 et Cie, Paris, pp. 3-7.

Honkela, T., Kaski, S., Lagus, K. and Kohonen, T. (1996), *Newsgroup Exploration with WEBSOM Method and Browsing Interface*, Report A32, Helsinki University of Technology, Laboratory of Computer and Information Science.

Honkela, T. (1997), "Learning to Understand - General Aspects of Using Self-Organizing Maps in Natural Language Processing" in Dubois, D. (ed.), *Computing Anticipatory Systems*, American Institute of Physics, Woodbury, New York, pp. 563-576.

Honkela, T. (2000a), "Adaptive and Holistic Knowledge Representations Us-

ing Self-Organizing Maps" in Shi, Z., Faltings, B. and Musen, M. (eds.), Proceedings of IIP'2000, International Conference on Intelligent Information Processing, International Federation on Information Processing, Beijing, China, pp. 81-86.

Honkela, T. (2000b), "Self-Organizing Maps in Symbol Processing" in Wermter, S. and Sun, R. (eds.), Hybrid Neural Systems, Springer, Heidelberg, pp. 348-362.

Honkela, T., Leinonen, T., Lonka, K. and Raike, A. (2000), "Self-Organizing Maps and Constructive Learning" in ICEUT Board (ed.), Proceedings of ICEUT'2000, International Conference on Educational Uses of Communication and Information Technologies, International Federation on Information Processing, Beijing, China, pp. 339-343.

Honkela, T., Hynnä, K.I. and Knuuttila, T. (2003), "Framework for Modeling Partial Conceptual Autonomy of Adaptive and Communicating Agents" in Alterman, R. and Kirsh, D. (eds.), Proceedings of Cognitive Science Conference 2003 (CD-ROM), Cognitive Science Society, Cincinnati.

Honkela, T. and Winter, J. (2003), Simulating Language Learning in Community of Agents Using Self-Organizing Maps, Report A71, Helsinki University of Technology, Publications in Computer and Information Science,

Hutchins, W.J. (1986), Machine Translation: Past, Present, Future. Ellis Horwood, Chichester.

Isabelle, P. (1987), "Machine translation at the TAUM group" in King, M. (ed.), Machine translation today: the state of the art, Edinburgh University Press, Edinburgh, pp. 247-277.

Jäppinen, H. and Ylilampi, M. (1986), "Associative Model of Morphological Analysis: An Empirical Inquiry", Computational Linguistics, vol. 12, no. 4, pp. 257-272.

Jäppinen, H., Honkela, T., Hyötyniemi, H. and Lehtola, A. (1988), "A Multilevel Natural Language Processing Model", Nordic Journal of Linguistics, vol. 11, pp. 69-87.

Karlsson, F. (1999), Finnish: An Essential Grammar, Routledge.

Kaski, S., Honkela, T., Lagus, K. and Kohonen, T. (1998), "WEBSOM -

self-organizing maps of document collections", *Neurocomputing*, vol. 21, pp. 101-117.

Kohonen, T. (1982), Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, vol. 43, pp. 59-69.

Kohonen, T. (1988), "The 'neural' phonetic typewriter", *Computer*, vol. 21, no. 3, pp. 11-22.

Kohonen, T. (1993), "Physiological interpretation of the self-organizing map algorithm", *Neural Networks*, vol. 6, pp. 895-905.

Kohonen, T. (1996), "Emergence of invariant-feature detectors in the adaptive-subspace self-organizing map", *Biological Cybernetics*, vol. 75, pp. 281-291.

Kohonen, T., Oja, E., Simula, O., Visa, A. and Kangas, J. (1996), "Engineering applications of the self-organizing map", *Proceedings of the IEEE*, vol. 84, no. 10, pp. 1358-1384.

Kohonen, T., Kaski, S., and Lappalainen, H. (1997), "Self-organized formation of various invariant-feature filters in the adaptive-subspace SOM", *Neural Computation*, vol. 9, pp. 1321-1344.

Kohonen, T. (2001), *Self-Organizing Maps*, 3rd edition, Springer.

Koskenniemi, K. (1983), *Two-level Morphology*, Ph.D. thesis, University of Helsinki.

Lagus, K., Honkela, T., Kaski, S. and Kohonen, T. (1999), "WEBSOM for Textual Data Mining", *Artificial Intelligence Review*, vol. 13, pp. 345-364.

Lehtola, A., Tenni, J., Bounsaythip, C. and Jaaranen, K. (1999), "WEBTRAN: A Controlled Language Machine Translation System for Building Multilingual Services on Internet" in *Proceedings of Machine Translation Summit VII*, Singapore, pp. 487 - 495.

Lin, X., Soergel, D. and Marchionini, G. (1991), "A self-organizing semantic map for information retrieval" in Bookstein, B., Chiaramella, Y., Salton, G. and Raghavan, V.V. (eds.), *Proceedings of 14th Annual International ACM/SIGIR Conference on Research and Development in Information Retrieval*, pp. 262-269.

- Maturana, H. and F. Varela (1980), *Autopoiesis and Cognition: The Realization of the Living*, D. Reidel Publishing Co., Dordrecht, Holland.
- Merkl, D. (1994), "Self-Organization of Software Libraries: An Artificial Neural Network Approach", PhD thesis, Institut für Angewandte Informatik und Informationssysteme, Universität Wien.
- Miikkulainen, R. (1993), *Subsymbolic Natural Language Processing: An Integrated Model of Scripts, Lexicon, and Memory*, MIT Press, Cambridge, MA.
- Nirenburg, S., Carbonell, J., Tomita, M. and Goodman, K. (1992), *Machine Translation: A Knowledge-Based Approach*, Morgan Kaufmann Publishers, San Mateo, CA.
- Ritter, H. and Kohonen, T. (1989), "Self-organizing semantic maps", *Biological Cybernetics*, vol. 61, no. 4, pp. 241-254.
- Rumelhart, D.E., Hinton, G.E., and Williams, R.J. (1986), "Learning internal representations by error propagation" in Rumelhart, D.E., McClelland, J.L. and the PDP Research Group (eds.), *Parallel Distributed Processing, Explorations in the Microstructure of Cognition. Volume 1: Foundations*, MIT Press, Cambridge, MA, pp. 318-362.
- Simula, O., Ahola, J., Alhoniemi, E., Himberg, J. and Vesanto, J. (1999), "Self-Organizing Map in Analysis of Large-Scale Industrial Systems" in Oja, E. and S. Kaski (eds.), *Kohonen Maps*, Elsevier, pp. 375-387.
- Steels, L. (1996), "Synthesising the origins of language and meaning using co-evolution, self-organisation and level formation" in Hurford, J., Knight, C. and Studdert-Kennedy, M. (eds.), *Evolution of Human Language*, Edinburgh University Press, Edinburgh, pp. 384-404.
- Valkonen, K., Jäppinen, H. and Lehtola, A. (1987), "Blackboard-based Dependency Parsing" in McDermott, J.P. (ed.), *Proceedings of IJCAI'87, International Joint Conference on Artificial Intelligence*, pp. 700-702.
- Visa, A. and Iivarinen, J. (1997), "Evolution and Evaluation of a Trainable Cloud Classifier", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 35, no. 5, pp. 1307-1315.
- Von der Malsburg, C. (1973), "Self-organization of orientation sensitive cells

in the striate cortex", *Kybernetik*, vol. 14, pp. 85-100.

Von Foerster, H. (1981a), "Notes on an epistemology for living things" in Von Foerster, H., *Observing Systems*, Intersystems Publications, pp. 257-271. Originally published in 1972 as BCL Report No 9.3., Urbana, Biological Computer Laboratory, University of Illinois.

Von Foerster, H. (1981b), "Perception of the future and the future of perception", in Von Foerster, H., *Observing Systems*, Intersystems Publications, pp. 189-204. Originally published in 1972 in *Instructional Science*, vol. 1, no. 1, Elsevier/North-Holland, pp. 31-43.

Von Foerster, H. (1981c), "Responsibilities of Competence" in Von Foerster, H., *Observing Systems*, Intersystems Publications, pp. 206-210.

Von Foerster, H. (1984), "Principles of Self-Organization in a Socio-Managerial Context" in Ulrich, H. and Probst, C.J.B. (eds.), *Self-Organization and Management of Social Systems*, Springer, Berlin, pp. 2-24.

Vuori, V., Laaksonen, J., Oja, E., and Kangas, J. (2001), "Experiments with adaptation strategies for a prototype-based recognition system for isolated handwritten characters", *International Journal on Document Analysis and Recognition*, vol. 3, pp. 150-159.

Winograd, T. (1972), *Understanding Natural Language*, Academic Press, New York.

Zadeh, L.A. (2002), "Precisiated Natural Language-Toward a Radical Enlargement of the Role of Natural Languages in Information Processing, Decision and Control" in Wang, L., Halgamuge, S.K. and Yao, X. (eds.), *Proceedings of FSDK'02, the 1st International Conference on Fuzzy Systems and Knowledge Discovery: Computational Intelligence for the E-Age*, Singapore, pp. 1-3.