TKK Reports in Information and Computer Science

Espoo 2009 TKK-ICS-R24

MODELING COMMUNITIES OF EXPERTS

Conceptual grounding of expertise

Timo Honkela, Nina Janasik, Krista Lagus, Tiina Lindh-Knuutila, Mika Pantzar and Juha Raitio

TEKNILLINEN KORKEAKOULU
TEKNISKA HÖGSKOLAN
HELSINKI UNIVERSITY OF TECHNOLOGY
TECHNISCHE UNIVERSITÄT HELSINKI
UNIVERSITE DE TECHNOLOGIE D'HELSINKI

TKK Reports in Information and Computer Science
Espoo 2009 TKK-ICS-R24

MODELING COMMUNITIES OF EXPERTS

Conceptual grounding of expertise

Timo Honkela, Nina Janasik, Krista Lagus, Tiina Lindh-Knuutila, Mika Pantzar and Juha Raitio

Helsinki University of Technology
Faculty of Information and Natural Sciences
Department of Information and Computer Science

Teknillinen korkeakoulu
Informaatio- ja luonnontieteiden tiedekunta
Tietojenkäsittelytieteen laitos

Distribution:
Helsinki University of Technology
Faculty of Information and Natural Sciences
Department of Information and Computer Science
P.O.Box 5400
FI-02015 TKK
FINLAND
URL: http://ics.tkk.fi
Tel. +358 9 470 01
Fax +358 9 470 23369
E-mail: series@ics.tkk.fi

ISBN 978-951-22-262-4 (Print)
ISBN 978-951-22-263-1 (Online)
ISSN 1797-5034 (Print)
ISSN 1797-5042 (Online)
URL: http://lib.tkk.fi/Reports/2009/isbn978951222631.pdf

TKK ICS
Espoo 2009

ABSTRACT: Finding ways in which communities of experts can benefit from each other is a question shared by the machine learning community and social sciences alike. Considerable research in machine learning methods has shown that communities of experts can provide consistently better classifications and decisions than single experts in various tasks and domains. Our aim is to extend the perspective on communities of experts to cover the wider context of socio-cognitive research. In particular, we discuss the socio-cognitive research on the formation and use of expertise in relation to the modeling of concept formation, integration and use in human and artificial agents. We present three case studies related to problem solving and decision making in environmental policy, medical care, and consumer research. We present a methodological framework for the computational modeling of the phenomena described above. A specific emphasis is on unsupervised statistical machine learning of heterogeneous conceptual spaces in multi-agent systems and on the application of such conceptual expert knowledge.

KEYWORDS: Expertise, computational modeling, machine learning, implicit knowledge, subjectivity, intersubjectivity, social simulation, Bayesian methods, self-organizing map

ACKNOWLEDGEMENT: We gratefully acknowledge the financial support for our research from Academy of Finland (TH, KL, TL-K, MP), TEKES - the Finnish Funding Agency for Technology and Innovation (TL-K, JR), and Helsinki University of Technology (NJ). TH, KL, TL-K, and JR are from the Adaptive Informatics Research Centre, Department of Information and Computer Science, Helsinki University of Technology, NJ is from the Department of Civil and Environmental Engineering, Helsinki University of Technology, and MP is from the Helsinki School of Economics (on leave from the National Consumer Research Centre). As a result of an upcoming university merger, the Helsinki University of Technology will form the School of Science and Technology and the Helsinki School of Economics will form the School of Economics in the new Aalto University from the beginning of 2010.

CONTENTS

1 INTRODUCTION

In cognitive science, artificial intelligence and machine learning research, it has been commonplace to concentrate on understanding and modeling the expertise of a single agent. In this report, we discuss some aspects of the expertise of an individual and describe why it is important to consider also the social level. We follow a socio-cognitive approach, i.e., an approach in which socio-cognitive processes are taken to be a complex and dynamic combination of coupled individual cognitive processes. These couplings give rise to the social level of abstraction.

We provide a summary of our earlier research related to this field and outline an analytical and methodological framework for the socio-cognitive modeling of communities of experts. In section 2, we report earlier research related to the ways of understanding and modeling individual expertise. We make a distinction between explicit and implicit knowledge of an expert. Methodologically, we provide a broad overview including discussion on rule-based representation, artificial neural networks, and Bayesian modeling. Related to the complexity of the underlying phenomena, we discuss the relation between holism and reductionism from an epistemological point of view. In section 3, we expand the view to the social level, concentrating on the notion of distributed expertise and ensemble learning. In section 4, we provide three different empirical perspectives on expertise. The three domains are natural resource use, health care, and consumer behavior. In these, we problematize the relation between "lay" and "expert" knowledge.

In section 5, we first describe how the subjective conceptual understanding underlying an expert's knowledge can be modeled without assuming an innate existence of a conceptual system. Next, we discuss how an intersubjectivity between subjective conceptual systems can be reached as an adaptive process. Finally, we outline three distinct ways in which conceptual structure can be integrated in processes of expert communication with the aim of solving problems. These three strategies are, in order of increasing complexity, a) clarifying naming conventions, b) visualizing differences in conceptual density, and c) providing augmenting data that mediates between the different conceptual systems.

Based on our analytical and methodological framework, it seems that which kind of strategy of conceptual integration that is most suitable in a specific collaborative process between experts depends on the nature of the confronted problems. It is important to note that the processes of conceptual interaction and integration that we describe in this report are not level-dependent, i.e., they remain the same irrespectively of whether they occur at, or across, the levels of individual, group or unit.

This report describes the area of modeling expertise in a multi- and interdisciplinary manner bringing together results from cognitive science, sociology, artificial and computational intelligence (esp. statistical machine learning research), and science and technology studies. We are in the process of developing the analytical and methodological framework into a useful theoretically viable tool for the analysis and facilitation of expert collaboration.

2 UNDERSTANDING AND MODELING INDIVIDUAL EXPERTISE

Next, we consider different computational models that have been used to represent individual expertise. In particular, we make a distinction between explicit representations (such as rule systems) and implicit representations (such as artificial neural networks). In addition to presenting various computational methods in this field, we relate these methods with empirical finding in the human sciences.

2.1 Individual expertise represented as rules

In the 1980's it was commonplace to view expertise as a collection of explicit rules. A wide variety of expert systems were built in order to codify the knowledge of the experts of some specific domain (consider ,e.g., [2, 99]. Typically, such an expert system consists of 1) a *knowledge base* encoded in some formalism, closely related to predicate logic, and 2) a *reasoning mechanism* to device inferencing that is based on the rules of the knowledge base. The reasoning mechanism is implemented as an inference engine. The rules in an expert system are usually production (if-then) rules. The reasoning mechanism can be based on forward chaining or backward chaining strategies. Prolog as a commonly used logic programming language [102] applies backward chaining which made it popular as an implementation platform for expert systems [71]. The underlying resolution mechanism [61] is implemented using Horn clauses. A Horn clause is a propositional clause (a disjunction of literals) with at most one positive literal. This kind of Horn clause, i.e.,

$$\neg p \vee \neg q \vee \cdots \vee \neg t \vee u$$

can be written equivalently in the form of an implication, i.e.,

$$(p \wedge q \wedge \cdots \wedge t) \Rightarrow u$$

The experiences from a large number of expert system development projects also highlighted some problems in representing expertise in the form of explicit rules. The first class of problems is quantitative. In many domains, the number of rules needed to represent the body of knowledge is large. The amount of work needed for the knowledge acquisition process for a typical domain could be several person years. For the experts that were interviewed to collect the expertise, it was also often difficult to express they knowledge in the form of explicit rules. Moreover, the rules provided by different experts could contradict with each other which is closer to the second class, i.e. qualitative problems. For instance, rules do not suit well for the representation of imprecise information. As an early solution, certainty values were associated with each rule. This was not a theoretically grounded approach and it has since been replaced by the use of Bayesian reasoning techniques. Bayesian inference uses an estimate of the degree of belief in a hypothesis before evidence has been observed and based on observed evidence, calculates an estimate of the degree of belief in the hypothesis. As a model of expertise, the Bayesian model is interesting as it allows to present the initial set of beliefs and a systematic update of the beliefs based on experience. Bayes' rule adjusts probabilities given new evidence as follows:

$$P(H|E) = \frac{P(E|H)\,P(H)}{P(E)}$$

In this well known equation, H represents a specific hypothesis. $P(H)$ is the prior probability of H before new evidence (E) became available. $P(E|H)$ is the conditional probability of seeing the evidence E if the hypothesis H happens to be true. $P(H|E)$ is the posterior probability of H given E. $P(E)$ is the marginal probability of E. In other words, $P(E)$ is the a priori probability of witnessing the new evidence E under all possible hypotheses. It can be calculated as the sum of the product of all probabilities of any complete set of mutually exclusive hypotheses and corresponding conditional probabilities:

$$P(E) = \sum P(E|H_i)P(H_i).$$

In the Bayesian framework, accumulating expertise can be viewed as a process of collecting an increasing number of rules with associated conditional probabilities. The framework can be used to model expertise at different levels of explicitness. Next, the explicitness-implicitness question is considered in some detail.

2.2 Implicit expertise of an individual

Empirical research on human knowing and experience has clearly shown that expertise is based on skills and knowledge that are difficult to represent explicitly in linguistic form. The term "tacit knowledge" is often used to refer to knowledge that is difficult to be transferred from one person to another by means of writing down or verbalizing it [78, 75]. Dijksterhuis et al. have recently shown that unconscious or intuitive decision making gives systematically better results than reliance on explicit or rational thinking in solving complex problems [19]. Evans points out that people reason in a probabilistic manner [23]. Moreover, reasoning is highly contextualized by relevant prior knowledge and belief. In the dual process theories of reasoning, a division is made between a heuristic system and an analytical system. The heuristic system has evolved early, it is shared with animals, it is rapid and parallel, it has high capacity and it is pragmatic [23]. These important aspects of knowledge and reasoning are set aside if they are only considered at the level of the analytical system that has evolved late and that provides those kinds of thinking tools that are necessary for, e.g., logical reasoning.

Considering the analytical level and explicit representation is not enough as knowledge is based on the underlying experiential domain. Gigerenzer describes this phenomenon in a very clear manner in his book "Gut Feelings - The Intelligence of the Unconscious" [32]. The section "Why good intuitions shouldn't be logical" is devoted to analyzing the classical reasoning task outlined by Tversky and Kahnemann to show how people tend to violate logical reasoning. In the task, people are asked which of the following alternatives is more probable: (A) Linda is a bank teller, or (B) Linda is a bank teller and active in feminist movement. The question is posed after explaining that Linda is 31 years old, outspoken, and very bright. As a student of philosophy she was deeply concerned with issues of discrimination and social justice and participated in antinuclear demonstrations. According to Tversky and Kahnemann's argumentation, the fact that people often choose B even though the choice A is right from the point of view of logic. A conjunction of two events cannot be more probable than only one of them. Gigerenzer criticizes this view to be *content-blind* because the content and the goals of thinking are ignored. He states that intelligence has to operate in an uncertain world, not in the artificial certainty of a logical system. In the Linda case, Gigerenzer showed that the majority of people understood the meaning of "probable" and "and" through their non-mathematical interpretation [32]. For instance, "probable" was interpreted as "conceivable", "plausible", "reasonable", and "typical". Moreover, it is clear that as a natural language expression the structure "(X and Y) or X" is easily understood as "(X and Y) or (X and not-Y')". The omission of "not-X" is an omission of an otherwise repeated structure which is a common phenomenon in language. From the point of view of expertise, the message is that problem solving in a case like the one described above is based on a model of the underlying phenomenon and an assessment of the intended task at hand.

In general, it seems that an individual's rationality is an adaptive tool that does not follow (only) the principles of symbolic logic or probability theory as such, but includes various "cognitive survival strategies", such as a collection of heuristics as pointed out by Gigerenzer and his colleagues [33]. For instance, the recognition heuristic states that "If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value with respect to the criterion." [33, 35] In the academic world, this kind of criterion can be used in decision making, e.g., in the following way. If one has found six articles related to some topic and one has time to read only one or two of them, the person is more likely to read, regardless of their individual quality, the articles coming from institutions like Cambridge, Harvard, or MIT, rather than Da Nang, Lappeenranta, or Tshwane.

2.3 Methods for implicit representation

The difference between explicit and implicit knowledge is usually defined by referring to language. If knowledge is represented as interpretable linguistic expressions, it is considered to be explicit, otherwise implicit. Computational intelligence methods such as neural networks and statistical machine learning have provided models of implicit (unconscious, intuitive) understanding. This link has occasionally been stated separately (see, e.g., [8]). These methods make it possible to model fine-grained many-to-many probabilistic relationships in knowledge and reasoning.

A feedforward multilayer perceptron can be seen to device a many-to-many mapping between two sets of variables. A multilayer perceptron is an artificial neural network model that maps sets of input data onto a set of output. It uses three or more layers of nodes with nonlinear activation functions. From the point of view of the present discussion, it is essential that every node in some layer is connected with every node in the following layer. This connectivity ensures that the network is able to perform as an arbitrary pattern classifier [68], a universal function approximator [45], or to be equivalent in power to a universal Turing machine [96]. Learning in a multilayer perceptron can take place with the backpropagation of errors algorithm. In this algorithm, connection weights w_{ji} between the nodes are changed based on the amount of error in the output compared to the expected result:

$$\Delta w_{ji}(n) = -\eta \frac{\partial \mathcal{E}(n)}{\partial v_j(n)} y_i(n)$$

In this equation, y_i is the output of the previous neuron and is the learning rate.

In such a network, the intermediate representation is implicit, even when the input and output layers contain variables with linguistic/explicit interpretation. The degree of implicitness grows, for instance, if the input layer contains variables that do not have (high-level) linguistic interpretation at a relevant level of abstraction. Of course, any representation based on distinct variables can be given a symbolic description ("the level of blue in the pixel in the row 121, column 73"). The term *relevant level of abstraction* can further be exemplified with recorded music. A digital recording of music consists of a huge number of individual bits but expressing each of these bits as a separately named entity does not make sense from the point of view of the musical level. A neural network model could be trained, for example, to serve as an expert of classical music in the sense that through an exposure to thousands of samples, the system would be able to recognize different composers with reasonably high accuracy. This recognition or classification skill would reside in the complex weight patterns of the neural classifier with limited means for making the reasons for each classification result explicit. Herbert Simon has expressed a related thought as follows: "The smartest people in the world do not generally look very intelligent when you give them a problem that is outside the domain of their vast experience."

Hyötyniemi has coined the term "neocybernetics" that builds on the idea of second-order cybernetics, as originally coined by von Foerster, as cybernetics of observing systems [27, 28]. Within neocybernetics, Hyötyniemi defines the basis for conceptual categories as follows: "The chunks (symbolic or subsymbolic memory elements) are patterns constructed of sparse-coded features, or 'relevant attractors' among data; inference is pattern matching in the data space." [48] He continues by stating, in general terms, that the process of shift from novice to expert can be explained in a numeric framework better than in a symbolic one. We agree with this statement, and in this report we provide a preliminary version of an analytical and methodological framework for representing and processing expert knowledge.

2.4 The issue of complexity

Weaver approached the concept of complexity by making a distinction between *disorganized complexity* and *organized complexity* [106]. According to Weaver, disorganized complexity results from the interactions between a very large number of parts in a particular system. Even though the interactions in this case would be more or less random, the properties of the system as a whole can be understood by using probability and statistical methods. Weaver used the number of balls on a billiards table as an illustrative example. The classical dynamics is well suited for analyzing the motion of a single ball. With increased difficulty, one can also analyze the motion of two or even of three balls on a billiard table. Referring to the analogy of analyzing gases, Weaver then asked to consider a large billiard table with millions of balls rolling over its surface, colliding with one another and with the side rails. In such a case, the methods of statistical mechanics are applicable [106].

From our point of view, a more interesting case among the two discussed above is the organized complexity, which incorporates non-random or correlated interaction between the parts of a system. In organized complexity, a coordinated system manifests *emergent* properties not dictated by individual parts.

A central question here is what is the relationship between holism and reductionism. We will consider this relationship from an epistemological point of view. Holism then refers to the idea that all the properties of a given system cannot be explained by its component parts alone[1]. The system as a whole is needed in the explanation of how the parts behave. Reductionism, as an opposite to holism, says that a complex system can be explained by reduction to its fundamental parts. It seems that both extremes are problematic. If holism is taken to the extreme, it would mean that we would need to analyze the connections between all elements at all levels of abstraction. On the other hand, a fully reductionist approach does not seem to be fruitful either. As an intermediate view, Simon has introduced the concept of near-decomposability [97]. As a property that seems to be shared by all multicelled organisms, near-decomposability refers to hierarchies of components, such that, at any level of the hierarchy, the rates of interaction within components at that level are much higher than the rates of interaction between different components [97]. From the evolutionary point of view, under the usual conditions of mutation and/or crossover and natural selection, nearly decomposable systems will increase in fitness, and therefore reproduce, faster than systems without this property.

An analogical structure to Simon's concept of near-decomposability can be discerned in many areas of science, technology and their methodologies. In clustering methods, the objective is to find collections of objects which are similar with each other and dissimilar to the objects belonging to other clusters. In software engineering, it is commonplace to use modular structures in which the data flow between the modules is restricted. In formal linguistics, compositionality is considered to be an important characteristic with the recognition that exceptions exist. As a related theme, natural language parsing is often conducted using context-free grammar formalisms even though context sensitivity is needed in many cases due to long distance dependencies. It is important to recognize that the "truth" lies somewhere between the two extremes as suggested by Simon's term near-decomposability.

2.5 Bayesian framework

From the methodological point of view, it may be useful to review some aspects related to Bayesian modeling. The main idea of this subsection is

[1]Ontologically, the question would rather be about determination than explanation.

to consider how complexity of representing knowledge is handled within a framework that is widely used in science. Let us first remind ourselves of the basics of the Naive Bayes method[2]. A naive Bayes classifier is a probabilistic classifier based on applying Bayes' theorem (discussed earlier in this report). Strong independence assumptions are in use which refers to the assumption that the presence (or absence) of a particular feature of a class is unrelated to the presence (or absence) of any other feature. Based on this assumption, the Naive Bayes approach can be called reductionist when considered from the point of view of the discussion above. In concrete terms, the Naive Bayes classifier is based on the following conditional model in which class variable C is conditional on several feature variables, $F_1 \dots F_n$.

$$p(C|F_1, \dots, F_n)$$

Using Bayes' theorem, one can write this into the following form.

$$p(C|F_1, \dots, F_n) = \frac{p(C)\, p(F_1, \dots, F_n|C)}{p(F_1, \dots, F_n)}.$$

The denominator does not depend on C and the values of the features F_i are given, so that the denominator is effectively constant. The numerator is equivalent to the joint probability model.

$$p(C, F_1, \dots, F_n)$$

Using repeated applications of the definition of conditional probability, his can be rewritten as follows.

$$
\begin{aligned}
& p(C, F_1, \dots, F_n) \\
& = p(C)\, p(F_1, \dots, F_n|C) \\
& = p(C)\, p(F_1|C)\, p(F_2, \dots, F_n|C, F_1) \\
& = p(C)\, p(F_1|C)\, p(F_2|C, F_1)\, p(F_3, \dots, F_n|C, F_1, F_2) \\
& = p(C)\, p(F_1|C)\, p(F_2|C, F_1)\, p(F_3|C, F_1, F_2)\, p(F_4, \dots, F_n|C, F_1, F_2, F_3) \\
& = p(C)\, p(F_1|C)\, p(F_2|C, F_1) \dots p(F_n|C, F_1, F_2, F_3, \dots, F_{n-1}) \dots
\end{aligned}
$$

The complexity of this representation shows the effect of the "holistic" view in the values of the features may depend on each other. Now if we assume "reductionistically" that each feature F_i is conditionally independent of every other feature F_j for $j \neq i$, the mathematical representation becomes much simpler.

$$p(F_i|C, F_j) = p(F_i|C)$$

The joint model can then be expressed as follows.

$$p(C, F_1, \dots, F_n) = p(C)\, p(F_1|C)\, p(F_2|C)\, p(F_3|C)\, \cdots = p(C) \prod_{i=1}^{n} p(F_i|C).$$

In general, the independence assumption of the Naive Bayes method may be too restrictive in many real world situations. Simulation-based Monte Carlo techniques are often used in connection with Bayesian methods to overcome some of these problems (see, e.g., [90, 104]).

[2]The mathematical presentation of the method follows to a large degree the Wikipedia article "Naive Bayes classifier" at $http://en.wikipedia.org/wiki/Naive\_Bayes\_classifier$, downloaded 4th of November, 2009.

2.6 No free lunch theorem and inductive bias

There are some inherent restrictions related to the processes of generalization in learning. Already Hume [47] pointed out that the direct experience is the basis of knowledge.

> There is no object, which implies the existence of any other if we consider these objects in themselves, and never look beyond the ideas which we form of them. Such an inference would amount to knowledge, and would imply the absolute contradiction and impossibility of conceiving any thing different. But as all distinct ideas are separable, it is evident there can be no impossibility of that kind. When we pass from a present impression to the idea of any object, we might possibly have separated the idea from the impression, and have substituted any other idea in its room.
>
> It is therefore by EXPERIENCE only, that we can infer the existence of one object from that of another. The nature of experience is this.

From the point of view of forming conceptions of the world, the following passage is particularly interesting [47].

> Even after the observation of the frequent or constant conjunction of objects, we have no reason to draw any inference concerning any object beyond those of which we have had experience.

In the context of machine learning, in particular supervised learning, this thinking is famously formulated in the so called "No free lunch" theorem [108] (see also [109]). Wolpert shows that even if we know the classification error rate obtained by two algorithms a_1 and a_2 for a particular data set d_1, if we then take another, unknown data set d_2 which has no overlap with d_1, we still know nothing about which of the algorithms will be better. This holds regardless of the learning algorithms, and Wolpert shows that in particular it holds when a_1 is the algorithm favored by cross-validation and algorithm a_2 is the opposite, the one performing worst in cross-validation. Another way to express the gist of this work is to say that averaged over all possible problems, any two algorithms are equally good.

A computational learning problem can be often viewed as a search for a good model from a space of candidate models. What makes learning algorithms different, is their so-called *inductive bias*. One way to explain inductive bias is to say that it is the policy using which the algorithm searches the space of possible models. Two algorithms with a different inductive bias will arrive in different areas of the search space with different probabilities. This means that if one algorithm is superior with certain kinds of data sets, the other algorithm must be superior with some other kinds of data sets: averaged over all possible data sets the algorithms are equally good.

In practice, most algorithms only evaluate a small subset of all possible candidate solutions. However, since often the search spaces are such that good and even optimal solutions are found scattered all over the search space, it is likely that after examining a quite small set of candidate solutions a rather good solution will be found by any algorithm. Restricting the search space of candidate models using prior information regarding, for example, the domain of application can be viewed as a useful approach.

For our purposes, the computational issues discussed above (as well as Hume's observations) seem to point out that the experience of a phenomenon is important. When the number of observations related to a phenomenon is high, there is less need for extensive generalizations. This may be considered to be in line with the idea that the formation of *human expertise* takes several years to develop. Furthermore, this development of expertise requires enough direct exposure to the phenomenon at hand.

3 UNDERSTANDING AND MODELING EXPERTISE IN NETWORKS OF INDIVIDUALS

The social level of expertise refers to competencies that arise from social interaction, knowledge sharing, and collective problem solving ([36]. Cognition and intelligent activity are not only individual processes but ones which rely on socio-culturally developed cognitive tools. These include physical and conceptual artifacts as well as socially distributed and shared processes of intelligent activity embedded in complex social and cultural environments [36]. Expertise at the social level is constituted in interaction between individuals, communities, and larger networks supported by cognitive artifacts.

At the socio-cultural level, humans create and share conceptual artifacts such as symbols, words and texts. These are used as mediators between different minds. In communicating and sharing knowledge, individuals have to make a transformation between their internal representation into an explicit representation to be communicated and vice versa, as Vygotsky pointed out already in the 1930s. The internalization and externalization processes take place as a continuous activity. In externalization, the internal view is externalized as explicit and shared representations.

In their book, Castellani and Hafferty provide a careful account of scientific work at the intersection of complexity science and sociology. They consider, among other things, such areas as computational sociology, sociocybernetics, and social network analysis. Citing, e.g., Buckley [11] and Luhmann [70], the authors state that studying society is the same as studying complexity. Methodologically, they list a number of subdisciplines and concepts of relevance for this area including systems science, cybernetics, dynamic systems theory, artificial intelligence, neural networking, self-organization, autopoiesis, agent-based modeling, genetic algorithms, and artificial life[3].

3.1 Distributed expertise and ensemble learning

In distributed AI and multi-agent system research the idea of using multiple experts to solve a problem in collaboration is common. The use of a com-

[3] A diagram illustrating these areas and their relationships is available at $http://www.personal.kent.edu/\tilde{}bcastel3/complex\_map.html$

munity of experts may be motivated by the fact that various agents master different parts of the domain in question or they apply different approaches (methods, parameterizations, initializations) the performance of which cannot be determined beforehand in different cases (see [93] for more information on multiagent systems).

A methodologically sophisticated approach in this area is ensemble learning. Ensemble methods have been applied in various forms and under various names to classification, regression and time series prediction. A non-exhaustive list of approaches for combining different models into a single model includes bagging, boosting, committees, mixture of experts, multi-agent systems for prediction, and classifier ensembles (see [38] for further references). The classical mixture of experts model generates classifier experts [52]. Their outputs are combined through a generalized linear rule. In general, ensemble learning is the process by which multiple models, such as classifiers or experts, are strategically generated and combined to solve a particular computational intelligence problem [79]. Dietterich mentions statistical, computational and representational reasons for using ensemble based systems [18]. There may not be adequate data available to properly represent the data distribution (statistical reason). There are many alternative models that can be used to solve the given problem (computational reason). Sometimes the chosen model does not include a sufficiently rich representation to facilitate adequate problem solving (representational reason). For instance, in a classification problem, the chosen model cannot properly represent the sought decision boundary. Using an ensemble of models, instead of choosing just one, and combining their outputs can reduce the risk of an unfortunate selection of a particularly poorly performing model [79]. In order for the problem solving process to be effective, the individual experts must exhibit some level of diversity.

3.2 Social simulation

Social simulation aims at exploring and understanding of social processes by means of computer simulation. Social simulation methods can be used to to support the objective of building a bridge between the qualitative and descriptive approaches used in the social sciences and the quantitative and formal approaches used in the natural sciences. Collections of agents and their interactions are simulated as complex non-linear systems, which are difficult to study in closed form with classical mathematical equation-based models. Social simulation research builds on the distributed AI and multi-agent system research with a specific interest of linking the two areas.

The research area of simulating social phenomena is growing steadily (see, e.g., [103]) but we do not aim at giving a comprehensive overview in this report. Rather, we next present and discuss two examples of recent research that is related to our topic.

Schwenk and Reimer have built a social simulation model to study the processes of social influence [92] that partly builds on the research on heuristics [33]. They examined the interaction of decision strategies and features of the communication network. Schwenk and Reimer's simulation model was contextualized by a scenario which they adapted from Lazega [65]. In this scenario, a group of lawyers who are partners in a law firm gather in a meeting in order to decide about topics concerning the firm, for instance, the branch of business in which the firm should further expand [92]. In the simulation model, the lawyers were represented by a set of 21 agents, each having a certain preference for a branch of business into which the firm should expand. In more detail, each agent l_i was associated with both a value d_i of a decision variable D, which contains the discrete decision alternatives, and a value w i of an individual status variable W. A directed graph G, described a network of directed communication channels c_{ji} between the agents L. Each agent l_i is assigned a decision procedure f out of a set of decision procedures F. This function f consisted of a contact rule r_c and a decision rule r_d and maps an agent's actual decision state d_{j_n} onto its subsequent state d_{j_n+1}. the dynamic evolution of the model was then based on the iterated and sequential call of this decision rule f.[92] The result of the simulation was such that the impact of the agents' decision strategies on the dynamics as well as the outcomes of the influence process depended on features of their social environment. This behavior particularly clear when the agents contacted all of the neighbors with whom they were connected [92]. From our point of view, an interesting extension of this work would be to combine the influence process modeling with some more specific consideration of the conceptual models of the agents.

Nishizaki, Katagiri and Oyama have developed an agent-based simulation model in which artificial adaptive agents have mechanisms of decision making and learning based on neural networks and genetic algorithms [74]. They compare the results of their simulation analysis with the ones of a mathematical model related to the potential occurrence of strikes in a labor market. One result stemming from the simulation model was that individuals behaved cooperatively and that the prisoners' dilemma could be escaped. The earlier model was based on rationality (individual utility maximization), whereas the agents in the social simulation behaved adaptively. Agents made decisions by trial and error, and they learned from experiences to make better decisions. [74]

3.3 Bayesian view on distributed expertise

In the Bayesian framework, the Bayes Optimal Classifier is an ensemble of all hypotheses in the hypothesis space. It can be expressed with the following equation[4]

$$y = \mathrm{argmax}_{c_j \in C} \sum_{h_i \in H} P(c_j|h_i)P(T|h_i)P(h_i)$$

In this equation, y is the predicted class, C is the set of all possible classes, H is the hypothesis space, and T is the training data. There are several reasons why this formulation as such is of mainly theoretical value. Many interesting hypothesis spaces are too large for the model. It is also non-trivial to compute an unbiased estimate of the probability of the training set given a hypothesis ($P(T|h_i)$. Finally, estimating the prior probability for each hypothesis,

[4]The mathematical presentation of this methods follows the Wikipedia article "Ensemble learning", $http://en.wikipedia.org/wiki/Ensemble\_learning$, downloaded 4th of November, 2009.

$P(h_i)$, is rarely feasible. In order to make the approach feasible, one solution is create an approximation of the above formula making simplifying assumptions. Rather than sampling all hypotheses, one approach is to use some Monte Carlo sampling technique such as Gibbs sampling. The relationship between the optimal solution and an approximation can be viewed from the point of view of dealing with the underlying complexity. In realistic contexts, the optimal solution is too complex ("holistic") and simplifying assumptions are needed to make the approach more tractable.

Within the Bayesian framework, there are also more specific approaches to consider the social level of expertise. For instance, Zhang has considered evolutionary learning at the population level [111]. According to him, individual animals/species increase or decrease their future probability of action choices based on the consequence of the currently selected action. Under Bayesianism, evidence is evaluated based on likelihood functions so that action probability is modified from a priori to a posteriori according to the Bayes rule. Viewed as hypothesis testing, an evolutionary/selectionist framework attributes evidence exclusively to the selected, focal hypothesis, whereas a Bayesian framework distributes across all hypotheses the support from a piece of evidence [111]. Zhang shows that when individuals modify their action choices based on the selectionists' approach, the learning population at the ensemble level evolves according to a Bayesian-like dynamics [111].

4 EMPIRICAL PERSPECTIVES ON EXPERTISE

In the following, we present three empirical contexts in which communication across knowledge boundaries or different forms of expertise takes place. All contexts challenge the traditional notion of expertise.

4.1 Natural resource users as experts

Experts, especially scientists and engineers, have always played a prominent role when taking policy action [16, 84]. Typically, a distinction has been made between *experts* in possession of *systematic knowledge*, and *lay persons* possessing only *contextual knowledge* [84]. During the 1950s and 1960s, it was inconceivable that decision making could travel any other direction than top-down [16]. Non-experts were seen as basing their *contextual* views on *values* and thus as incapable of understanding knowledge or of using it in a meaningful way [84]. As an academic movement, this way of thinking began to erode towards the end of the 1960s, being more or less superseded by the late 1970s [16]. Part of the erosion of the divide between lay and expert knowledge was the criticism from the perspective of social constructivism, that this *positivistic* stance fails to account for all the aspects of social life that go into the building of the scientific practice in the first place [16, 50]. Over and above this major twist, the criticism has taken a number of specific forms. For instance, Collins and Evans advocate the extension of *technical* expertise to include also *uncertified*, experience-based expertise [16]. Mitchell et al. propose that the technical expert be seen as an *honest broker* between various conflicting demands [72].

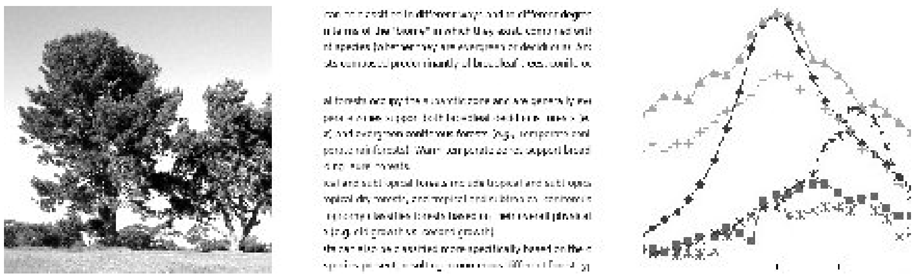

Figure 1: An illustration of different sources of knowing i.e., direct experience, written information, and numerical information.

Irwin argues that *scientific* and *contextual* knowledge need not necessarily conflict, since contextual knowledge may be constructed in scientific terms and since scientific knowledge can be seen as one kind of contextual knowledge [50]. According to Harrison et al., *local knowledge* can be more efficient in local contexts since it is specific to and historically embedded in this context [37]. Ravetz and others in the *post-normal science* camp argue that researchers need to accept other quality criteria than the scientific ones as legitimate in policy and decision making processes [82]. Post-normal science refers to the scientific inquiry when facts are uncertain, values are in dispute, stakes are high, and decisions are urgent. The debate over global warming can be seen as such an issue with other similar, long-term questions in which the scientific community possesses limited amount of information in relation to the importance of the questions at hand [30].

Others, however, criticize the view that scientific and contextual knowledge can, in the last instance, be made compatible. This can be exemplified within the context of natural resource users and managers, such as sheep farmers [110] or herders [46]. In his famous study of the ways in which scientists and sheep farmers deal with radioactive pollution in North Cumbria pastures, Wynne highlights the ways in which the scientific perspective of the radiologists and the practice-based perspective of the farmers differed not only with regard to what counts as knowledge in the first place (prediction and control vs. adaptability and flexibility in the context of action), but also, as inextricably linked to this, with regard to what is to be viewed as desirable from the point of view of both practical adequacy and morality (i.e. the latter as opposed to the former) [110]. Otherwise put, these two kinds of *knowledge practices* form *bundles* in which fact and value are inextricably intertwined, and there is no easy way of translating the one into the other. Recent research from the field of natural resource management seems to corroborate this view (see, e.g., [51, 46]). Studying the processes by which various kinds of not necessarily compatible expert perspectives are contextually generated thus yields the view that expertise is a phenomenon non-reducible to *technical* expertise, however broadly defined; that there needs to be a sensitivity to the ways in which public issues are framed and given meaning so that one kind of knowledge commitment is not allowed to hegemonize policy and decision making; and that there needs to be a sensitivity to the way in which all kinds of expertise are interwoven with deep issues of identity and *belonging* [110, 54, 55, 46]. The methodological framework constructed in this report is especially important for capturing, in a form easily understood by "tradi-

Figure 2: The MedIEQ project, funded by the EU Commission under the Public Health programme, has been conducting research related to the interface between layperson and expert knowledge in medicine. The main objective of the project was to develop methods and tools to help in ensuring the quality of health web sites (see *http : //www.medieq.org/* for more information.

tional" experts, the kind of expertise formed through centuries of close and committed (inter)action in and with complex ecosystems.

4.2 Patients as experts

The divide between experts and lay persons has been studied also within the context of the relationships between doctors and patients [15, 29, 22]. The study by Epstein tells the story of how AIDS patients in the US have managed to slowly establish themselves as credible participants in the construction of biomedical knowledge, as capable of providing insight and experience of their own condition [22]. Part of this has been learning to *speak the language* of biomedicine [22]. It is thus an exemplification of a successful attempt at creating a "new form of (somatic) subjectivity through collective activity" [87]. Experiences of this kind show the growing need to understand more closely the different types of expertise, and the types of collective conceptual processes and information sharing that can in principle take place, if the circumstances are favorable. Other examples can be found in the patient collectives that have formed around breast cancer and Lyme disease [22].

In the medical domain, the experts can thus be seen to have both different input knowledge bases (data sets, feature spaces and *top-down* influences) and different capabilities for action. The knowledge base of the doctor obviously is heavily affected by the medical education and research. A general practitioner's concept space is rather evenly distributed in the medical domain, whereas the conceptual space of a specialist doctor contains finer conceptual distinctions in a specific domain. The patient's data, especially in the case of a chronic illness or condition, is concentrated on personal experience and information that is most relevant to this particular case, gathered through the kind of active learning that is based on attentional guidance. Given the current medical resources in the internet, the active patient may over time become a very specialized expert also regarding medical research about the condition. The different experts may also have different, complementary capabilities for action. One (the doctor) is able to issue drug treatment, medical tests or surgery. The other (the patient) can start changes in life habits regarding exercise, nutrition, sleep or other everyday practices. Based on sharing expertise, the patient can also begin to collect further experiential data, paying attention to certain kinds of symptoms and phenomena, and monitor the daily effects of drugs and treatments for well-being. This is in essence what happened in the US examples described above. Unfortunately, however, not all circumstances are of this favorable kind. Bowker and Star [7] and Berg and Bowker [3] discuss the incompatibility between the medical categories and standards used by doctors and the experiences of patients. Discussing hormone replacement therapies in the UK, Roberts shows how menopausal women facing traditional medical expertise "struggle to articulate alternative knowledge claims or to pose potentially undermining questions" [87]. In many cases women *walk away* from clinical encounters with prescriptions for drugs they did not at all desire [87]. The knowledge gathered by both experts may be instrumental in deciding the best actions for each expert. If the knowledge sharing is successful, this may lead to better decisions for both. The methodological framework constructed in this report is of relevance in all communication between medical personnel and patients, but we consider it to be of especially high importance in circumstances characterized by Epstein as *less favorable* [22]. In particular, it can be used to capture the expertise of patients that have been characterized by Elbaz as "lay lay", i.e., as patient who actively participate in their own care but who have made the decision not to *learn to speak the language* of biomedicine [21, 22].

4.3 Consumers as experts

During the past few decades there has been a clear shift in consumer and innovation research from the viewpoint of the recipient of technology (the *Edisonian* technological and economic determinism) towards that of an active consumer. In this line of thinking, the consumers are no longer seen as mere recipients of technology, but even so, it is still customary to approach them via groups of elite consumers. The pioneers of technology (e.g., *lead users*, *pro-am consumers*, *users as producers*, or *innovative consumers*) are typically pictured as resourceful young men (see Fig. 1). In the early 21st century, this elitist viewpoint has been countered by a *social-movement* one: the masses make the movement. A new distribution technology (Web 2.0) and a new kind of willingness to follow the wisdom of the masses are at the core of Consumer 2.0, *Wikinomics* and *Democratizing of Innovations* ([39, 85]).

The distinction above between experts' and laypersons' expertise refigures in the context of innovation research, where the view of a sharp difference between core (professional lead users) and periphery (consumers) has recently been questioned also within the context of a general critique of *diffusionism* [4]. Diffusionism can be described as the belief that changes are produced by diffusion rather than by independent invention and that certain places are permanent centers of innovations. Diffusionism is a large and complex doctrine that has influenced many disciplines and countless arguments over 150 years [4]. Blaut questions the view according to which the core (professional lead users) is typically contrasted with the periphery (consumers). The terms used to denote this distinction are numerous: inventiveness vs. imitativeness, rationality vs. irrationality, intellect vs. emotion, abstract thought

Figure 3: The movie Star Wreck: In the Pirkinning is an example of user generated content. It is a full-length science fiction parody that took several years to make by five people in a two-room flat with a small budget and the support of a few hundred fans and dozens of acquaintances. The film is available online for free download and viewing under a Creative Commons license (see $http://www.starwreck.com/$ for more information. Even though no professional movie makers were involved, the movie has been estimated to be one of the most viewed movies ever produced in Finland.

vs. concrete thought, discipline vs. spontaneity, adult vs. child, and sane vs. insane [4]. In contrast to this view, we suggest that the difference between early adopters and late-comers is related to the fact that lead users have different elements of practices (skills, material objects, ideas and time resources) as compared to latecomers. There seems to exist social roles (and people) which act as *practice junctions* where elements (skills, ideas and objects) and practices cluster in a most dense way.

Today, ideas and material objects travel faster than ever (and our culture gives less and less universal cues how to behave correctly). Possibly it is through practice junctions that new products and practices get circulated. Lead users integrate differently practices, and as a result different practice constellations, e.g., TV dinners, emerge. Practice constellations are not only simple responses to existing needs [77]. They might have hidden transforming potential. It is especially new practices (elements and their combinations) that carry the seeds of emerging challenges and change. Possibly systems of practices are open to radical re-orientations only in their early stages (e.g., the use of the Internet in the mid 90's). For instance, the first digital item (and related practices) in a household may have been the kick-start for a whole new ecological system, where old species of analog technology become slowly replaced by other digital things, forming a community of interoperable goods and practices (e.g., digital photography [95]).

Bowden and Corkindale have demonstrated that occasionally it is experts and heavy users that are the most conservative when it comes to adopting new practices and adapting to radical innovations [6]. In general, however, it seems that it is often young people, information workers and those with much capital (cultural, economic) who are kinds of *settlements* for various new practices to attach. These people's life consists of number of small decisions which varies from one day to another. To overcome the complexity of stimuli those being in dense junctions (*practice squeeze*) have to negotiate and renegotiate their temporal order. This could result in feelings of inadequacy and hurriedness. An opposite social position is that of missing and lacking practices, inactivity and possibly feelings of loneliness. Seen this way individualism, adhocracy and feelings of hurriedness or loneliness do not arise from individual psychologies or aspirations but rather from different social positions in an array of daily practices. On the other hand, it is peripheral players who are more likely to adopt riskier innovations, because they have less at stake [88, 69]. In general, in complicated technology a proper lifestyle is critical for domestication [5, 20, 17].

Kotro has conducted research on hobbyist knowing, i.e. on how an essential part of product developers' knowledge stemmed from their free time hobbies [60]. The developers of sports measurement devices obtained grounded and embodied knowledge through their activities in scuba diving, mountain climbing, etc. [60]. What then makes being old (or young) so special in practice terms? Why do rich people seem to be more radical in one context and more conservative in another context? A hypothesis, based on views above, is that the difference between early adopters and late-comers is not related to psychological attributes but to the attributes of networks of practices and practitioners. First, lead users have different elements of practices (skills, material objects, ideas and time resources) compared to late comers. Second, lead users integrate practices differently, and as a result different practice constellations emerge. Third, as a result of successful integrations the very elements of practices get transformed (cultivation, learning) [77]. When discussing consumer expertise, we focus on concepts as part of the composite of *practice* understood as bundles of ideas, skills and objects. What is experienced and conceptualized as a *need* is at least partly an outcome of a learning process. We also argue that the methodological framework that we develop is of relevance for all three empirical contexts highlighted in this section. The same conditions for concept formation, integration and use across traditional boundaries of expertise apply to all of them, as does the need for processes of translation.

5 CONCEPTUAL VIEW ON EXPERT COMMUNITIES

In the following, we provide a theoretical and methodological framework for modeling concept formation and conceptual integration in expert communities. We also present some illustrative examples of the computational implementation.

5.1 Self-organizing conceptual spaces

Gärdenfors distinguishes between three cognitive levels of representation [31]. The most abstract level is the symbolic level, at which the information is represented in terms of symbols that can be manipulated without taking into account their meaning. The least abstract level is the subconceptual representation. Concepts are explicitly modeled at the mediating level of the conceptual representation.

A conceptual space is built upon geometrical structures based on a number of quality dimensions. Concepts are not independent of each other but can be structured into domains, e.g., concepts for colors in one domain, spatial concepts in the other domain. Fig. 4 shows an example of a conceptual space consisting of two quality dimensions, and two different ways (A and B) of dividing the space into concepts. A conceptual category c is seen as a convex region in a n-dimensional conceptual space C_i. The concepts are learned from a limited number of examples x and by generalizing from them. The similarity λ of two objects can be defined as a distance between their representation points in the conceptual space ($\lambda : C \times C \Rightarrow R$). The similarity measure can then be used for e.g. categorization[5]. A perceived item belongs to the category whose prototype is the nearest to the mapping of the item in conceptual space. In general, the theory of conceptual spaces proposes a medium to get from the continuous space of sensory information to a higher conceptual level, where concepts could be associated to discrete symbols. This mapping is a dynamic process. Gärdenfors [31] has proposed that for example multi-dimensional scaling (MDS) and self-organizing maps (SOM) [59] can be used to model this mapping process. The simplest connection between the SOM and conceptual spaces is to consider each prototype or model vector m in a SOM as an emerged conceptual category c. We can recall that the standard update rule of the SOM algorithm is as shown in Eq. 1 below. The learning phase, i.e. the iterative processing of the inputs $x(t)$ leads into an adaptation of the model vectors into an organized map. The function $h_{ci}(t)$ has a central role as it acts as a neighborhood function. In other words, h_{ci} is a smoothing kernel defined over the lattice of points [59].

$$m_i(t+1) = m_i(t) + h_{ci}(t)[x(t) - m_i(t)]$$

Research on conceptual modeling using the SOM includes the works by Ritter and Kohonen, Honkela et al., Lagus et al. and Raitio et al. [86, 42, 62, 81]. These results are discussed next in some detail. Related research has also been conducted, e.g., to study the initial conceptions of philosophy students [89] and the relation between the religiousness and counterintuitiveness of statements [80].

With a simple grammar, Ritter and Kohonen artificially generated syntactically correct and meaningful short phrases which were used as the input to the SOM [86]. The result was a map in which nouns, verbs and adverbs were clearly discernible. In this article, Kohonen discusses, among other things, the relationship between evolutionary and individual basis of concepts. He reminds us that at the time when the genetic predisposition of linguistic elements was suggested, there was no mechanism known that would have explained the origin of abstractions in neural information processing other than evolution. Kohonen further states that the "neural network" models are able to derive internal representations of categories from the mutual relations and roles of the primary signal or data elements themselves. This kind of emergence was demonstrated with two simple experiments based on the use of the SOM [86]. Regardless of the evidence that shows that linguistic abstractions can be learned with such models, there are researchers such as Chomsky who hypothesize that children have an innate knowledge of the basic grammatical structure [14]. It is further claimed that this structure would be common to all human languages. This innate knowledge is usually referred to as "universal grammar". Chomsky focuses on syntactic phenomena in language and, as

[5]It is to be noted that the concept of similary is closely related to the concept of analogy, e.g., in the research of Hofstadter [40].

such, the debate could be seen as mostly irrelevant from the point of view of modeling expertise – unless expertise in producing well formed expressions is in question. However, there are also prominent researchers such as Fodor whose position can be described as extreme or radical concept nativism. According to this view, virtually all lexical concepts are innate [24, 26]. This view has been analyzed critically by many, among whom Laurence and Margolis can be mentioned as a good example [64]. From the point of view of our argumentation, Fodor's radical concept nativism would mean that the conceptual grounding of expertise is straightforward and the differences in expertise would rather exist at the level of propositional knowledge. We reject this view as overly simplistic as it seems to overlook a large amount of empirical evidence. The problems related to polysemy or ambiguity can be mentioned as one example [1].

Honkela, Pulkki and Kohonen continued the work by Ritter and Kohonen by refining the method and by experimenting with segments of text from a natural corpus as an input for the self-organizing map [42]. The input for the map was the English translation of Grimm fairy tales. In the resulting map, in which 150 most frequent words of the tales were included, the verbs formed an area of their own in the top of the map whereas the nouns could be found in the opposite corner. The modal verbs were in one area. More semantically oriented ordering could also be found. For instance, the inanimate and animate nouns formed separate clusters. An important consideration is that in the experiments the input for the SOM did not contain any predetermined classifications. The results indicate that the text input as such, with the statistical properties of the contextual relations, is sufficient for automatic creation of meaningful implicit categories. The learning process gives rise to emergent categories. This kind of example stands in sharp contrast with the poverty of the stimulus argument often mentioned by the proponents of linguistic nativism. Argumentation like the one provided by Gold, who showed that any formal language which has hierarchical structure capable of infinite recursion is unlearnable from positive evidence alone [34], is not convincing because natural language does not exhibit infinite recursion. Actually, Karlsson has been able to show in a convincing manner that there are certain limits to the recursiveness of natural languages, in particular concerning center embedding [56]. Moreover, it has been shown that unsupervised systems that learn syntactic structures can be successfully deviced (see, e.g., [98]). (For detailed information on unsupervised learning, see [76].)

Lagus, Airola and Creutz analyzed the use of Finnish verbs to uncover possible conceptual spaces, and to study semantic similarities of verbs in actual language use. They examined the kinds of semantic or conceptual ordering qualities that appear to affect the distribution of features in the immediate context of a verb. With the unsupervised analysis based on the self-organizing map, the authors were able to find emergent categories like manipulative actions in human relationships, start of action (with focus on will or intention), communication (especially with positive emotional information), and aggressive or destructive use of power. This result further strengthens the argumentation presented above as no predefined categories were in use, i.e., only statistical co(n)textual information was present in the input.

Fodor and Lepore have presented an argument that connectionist theory

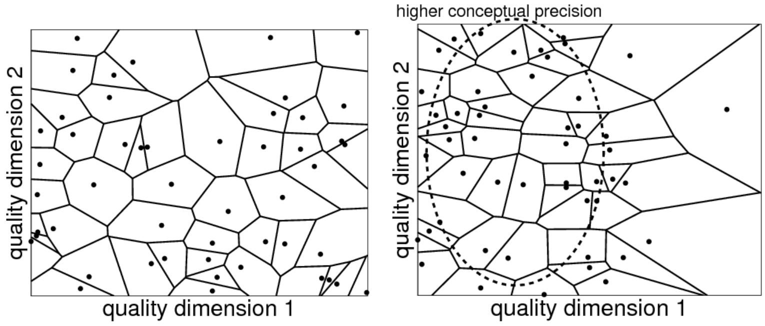

Figure 4: Illustration of differing conceptual densities of two agents having a 2-dimensional quality domain. Points mark the locations of the prototypes of concepts. Lines divide the concepts according to Voronoi tessellation. Both agents can discriminate an equal number of concepts, but abilities of the agent B are more focused on the left half of the quality dimension 1, whereas agent A represents the whole space with rather equal precision.

of mind cannot give a satisfactory account of different individuals being in the same mental state [25]. They claim that the identity of content follows from the identity of networks, but this condition will never be satisfied in practice. Raitio et al. have challenged this position by developing a methodology for comparing the similarity of representations in connectionist networks, and for examining the possibilities of exploiting it for comparing emergent representations in unsupervised learning networks [81]. This kind of comparison between different but with some respect essentially similar networks, points towards the consideration of subjective conceptual spaces and intersubjective mappings between the subjective spaces. These themes will be covered next.

5.2 Subjective conceptual spaces

Two persons may have very different conceptual density related to a topic under consideration. For instance, in Fig. 2 person A has a rather evenly distributed conceptual division of the space, whereas person B has a more fine-grained conceptual division on the left side of the conceptual space, but has lower precision on the right side of the space.

When language games are included in the simulation model, it resulted in a simple language emerging in a population of communicating autonomous agents [66]. In the population, each agent first learned a conceptual model of the world, in solitary interaction with perceptual data from the world. As a result, each agent obtained a somewhat different conceptual representation (a schematic illustration of the kinds of differences that can arise is shown in Fig. 2). Later, common names for the previously learned concepts were learned in communication with another agent.

5.3 Intersubjectivity in conceptual spaces

If some agents speak the *same language*, many of the symbols and the associated concepts in their vocabularies are the same. A subjective conceptual space emerges through an individual self-organization process. The input for the agents consists of perceptions of the environment, and expressions communicated by other agents. The subjectivity of the conceptual space of an individual is a matter of degree. The conceptual spaces of two individual agents may be more or less different. The convergence of conceptual spaces stem from two sources: similarities between the individual experiences (as direct perceptions of the environment) and communication situations (mutual communication or exposure to the same linguistic/cultural influences such as upbringing and education, and artifacts such as newspapers, books, etc.). In a similar manner, the divergence among conceptual spaces of agents is caused by differences in the personal experiences/perceptions and differences in the exposure to linguistic/cultural influences and artifacts.

The basic approach how autonomous agents could learn to communicate and form an internal model of the environment applying the self-organizing map algorithm was introduced, in a simple form, in [41]. The model has been later substantially refined in [43, 66, 44].

5.4 Modeling conceptually heterogeneous experts

When a community of conceptually heterogeneous human experts collaborate in order to solve challenging problems, for instance, in the environmental, health or consumer domains, they are likely to encounter a number of knowledge-related challenges [9]. Some of these challenges stem from differences in the conceptual systems of the individual experts. These kinds of situations call for means of highlighting the conceptual differences and resolving the resulting communication blocks. We present three strategies for this. These three strategies are, in order of increasing complexity, a) clarifying naming conventions, b) visualizing differences in conceptual density, and c) providing augmenting data that mediates between the different conceptual systems.

The computational models mentioned in Section 3 describe different principled ways in which artificial communities of experts can come to exist, and how their outputs can be combined.

In this section, we will look at the modeling of experts from a point of view that sheds light on the challenges of human problem solving that have their origin in conceptual differences. In particular, we will consider this issue from the point of view of the conceptual spaces framework described earlier.

Communication across borders of expertise in collaborative problem solving efforts can, in principle, be achieved in two ways: (1) by bringing forth a combination of the opinions of the experts by, e.g., voting, or (2) by a more involved sharing or integration of expertise and experience at the conceptual level. A particular form of sharing expertise is sharing prototypes. This refers to a process in which an expert communicates prototypical cases to the other expert. In the methodological context of the self-organizing map and other prototype-based conceptual models, prototype sharing means transmitting a collection of model vectors m_i.

Let us consider the features (essentially quality dimensions, see [31]) that span the conceptual space, data set (experience) used by an individual expert in learning the structure of its conceptual space, and the naming of concepts.

These three elements give rise to a typology of conceptual differences among experts. In the following, we present these different categories as well as the basic approaches for dealing with problems related to each category.

a) In the simplest case, the quality dimension space and data set are (nearly) equivalent for both agents. Only concept naming differs among different agents. An agent has an individual mapping function that maps each symbol to the conceptual space of the agent. In a classical simulation of this kind, a number of robots with cameras learned to name visual objects in a similar manner (see [101] and [107] as a philosophical background). An active research in language games and language evolution has since emerged (see e.g. [105, 66, 44]). Chen has presented a specific solution to the vocabulary problem among humans based on clustering [13]. Irwin's view that contextual knowledge may ultimately be constructed in scientific terms might be rooted in the view that differences in perspective are mainly a matter of concept naming [50]. This view might also figure in the background of much traditional or *standard* thinking in the domains of medicine and innovation.

b) As a step towards increased complexity, one may consider the situation in which the feature space is equivalent, but data set per expert varies. One expert has denser data from one part of the concept space, the other for another part (see Fig. 2). An obvious approach for efficient decision making is to use the expertise of those agents whose conceptual mapping is densest with regard to the problem at hand. However, in many cases, problem solving requires combination of many elements e.g. as solutions of subproblems. In those cases, each element can be dealt with by the expert with the densest conceptual mapping regarding a particular subproblem. Collins and Evans' advocation of the extension of *technical* expertise to include also *uncertified*, experience-based expertise might be rooted in the view that there exists a multitude of dense data sets, some of which are officially credentialized while others are not [16]. This view might also be behind calls for taking the views and experiences of patients more seriously, as well as behind recent calls to integrate the perspective of the user at an earlier stage in the innovation process than is often the case (e.g., [49]).

c) Finally, consider the most challenging case where neither the quality dimension space nor the data set are the same for both agents. Fig. 1 depicts a simple case in which the quality dimension spaces are different, therefore offering different viewpoints of the same *data sample* to the agents. In this case, a process of data augmenting can take place: if a subset of data samples known to both can be found (for example, boundary objects known across disciplines, or in terms of medicine, a particular patient's case), each agent can bring forth their particular knowledge (i.e., values of quality dimensions known only to them) regarding that case. Furthermore, in addition to collaborating in solving the present problem, both agents also have the opportunity to learn from each other: to augment their own representation with the new data offered by the other expert. Obtaining augmented information regarding several data samples will lead to the emergence of new, albeit rudimentary quality dimensions, and allow easier communication in future encounters. As an example, mutual data augmentation can take place between doctors of different specialization, doctors and patients, or between doctors and nurses, who consider simultaneously the same patient case. In optimal circumstances, this may eventually lead to better expertise of both. However, this requires that the doctor also trusts the patient, and is willing to learn and store the experiential data communicated by the patient. Essentially the same preconditions for and constraints to the process of data augmenting apply in the contexts of environmental policy and innovation.

6 CONCLUSIONS AND DISCUSSION

In this report, we have outlined an analytical and methodological framework for the socio-cognitive modeling of communities of experts. We have adopted a socio-cognitive approach, i.e., an approach in which socio-cognitive processes are taken to be a complex and dynamic combination of coupled individual cognitive processes, which then give rise to the social level of abstraction. In Section 2, we provided a broad overview including discussion on rule-based representation, artificial neural networks, and Bayesian modeling. In section 3, we expanded the view to the social level, concentrating on the notion of distributed expertise and ensemble learning. In section 4, we provided three different empirical perspectives on expertise, and focused on the problematic relation between between "lay" and "expert" knowledge. In all three empirical contexts this kind of data was profoundly grounded in real world experience: in the case of the Cumbrian sheep farmers, deep knowledge of the ecosystem obtained by centuries of interaction with it; in the case of patients, their own bodily experiences of what circumstances affect their condition in which way; in the case of innovations, experiences regarding mountains and deep sea regions.

Finally, in Section 5 we outlined three distinct ways in which conceptual structure can be integrated in processes of expert communication with the aim of solving problems. These three strategies are, in order of increasing complexity, a) clarifying naming conventions, b) visualizing differences in conceptual density, and c) providing augmenting data that mediates between the different conceptual systems. Based on this analytical and methodological framework, it seems that which kind of strategy of conceptual integration that is most suitable in a specific collaborative process between experts depends on the nature of the confronted problems. For example, a more well defined problem can be approached with the second (b) of these strategies, while a more ill defined problem requires strategies along the lines of (c). However, problem structures need to be seen as a continuum from a well defined to indefinite. Moreover, one needs to distinguish between different dimensions of a problem definition [63, 58].

The observation by Star and Griesemer on the significance of so-called boundary objects [100] provides an interesting point of access to these issues. A boundary object may provide a means for co-ordinating communication and action in situations in which the experts have different perspectives on some problem. From the point of view of the main theses of this report, we suggest that the nature and use of a boundary object be analyzed in more detail. Addressing these issues is part of our future work.

It is also important to note that the processes of conceptual interaction and integration that we describe in this report are not level dependent, i.e., they remain the same irrespectively of whether they occur at, or across, the levels of individual, group or unit. A central implication for real world situations across levels is that favorable circumstances for efficient problem solving in multi-expert situations include a process in which the participants become conscious of each others' type and domain of expertise. In particular, each expert needs to trust the data offered by the other experts. Previous research on, e.g., social capital [73] has shown that in order to reach such conditions, mutual respect and trust among the experts, and sufficient and open communication are needed.

Summarizing, by referring to conceptual spaces theory and the implementation of it using self-organizing neural networks we have provided an analytical and methodological framework for conceptual grounding of different kinds of expertise. The methodological framework presented in this report makes it possible to become conscious of important aspects of the expertise of other experts that otherwise easily gets marginalized out of the negotiation processes described in the three empirical contexts, as evidenced across all three domains (e.g., [60, 87, 110]. Most importantly, the methodology renders visible this other kind of expertise in a form that is more easily accessible to traditional experts. Thus, by enabling the identification and possibly prediction of conceptual overlaps and differences between various expert perspectives this analytical and methodological framework of modeling expert communities can significantly help reduce the cognitive barriers (for this and other kinds of barriers, see [10]) in expert communication in real life negotiation situations across epistemic boundaries.

REFERENCES

[1] K. Bach (2000) Review of Concepts: Where Cognitive Science Went Wrong. *San Diego Review*.

[2] A. Barr & E.A. Feigenbaum (1981) *The Handbook of Artificial Intelligence*. Volume 1, London: Pitman Books.

[3] M. Berg & G. Bowker (1997) The Multiple Bodies of the Medical Record. *Sociological Quarterly*, 38(3): 513-537.

[4] J. Blaut (1987) Diffusionism: A Uniformitarian critique. *Annals of the Association of American Geographers*, 77(1): 30-47.

[5] Bouwhuis, D. (2000) Parts of life: configuring equipment to individual lifestyle. Ergonomics, 43:908-919.

[6] J. Bowden & D. Corkindale, D. (2005) Identifying the initial target consumer for innovations: an integrative approach. *Marketing Intelligence & Planning*, 23(6): 562-573.

[7] G.C. Bowker & S.L. Star (1999) *Sorting Things Out: Classification and its Consequences*. Cambridgde, MA: The MIT Press.

[8] S. Brunak & B. Lautrup (1990) *Neural networks: computers with intuition*. World Scientific Publishing Co., Inc., River Edge, NJ.

[9] H. Bruun, J. Hukkinen, K. Huutoniemi, J. Thompson-Klein (2005) *Promoting Interdisciplinary Research: The Case of the Academy of Finland*. Helsinki: Publications of the Academy of Finland 8/05.

[10] H. Bruun, R. Langlais & N. Janasik (2005) Knowledge networking: A conceptual framework and typology. *VEST*, 18(3-4): 73-104.

[11] W. Buckley (1998) *Society–a complex adaptive system: essays in social theory*. Taylor & Francis.

[12] B. Castellani & F.W. Hafferty (2009) *Sociology and Complexity Science: A New Field of Inquiry*. Springer.

[13] H. Chen (1994) Collaborative systems: solving the vocabulary problem. *IEEE Computer*, 27(5):58-66.

[14] N. Chomsky (1975) *Reflections on Language*. New York: Pantheon Books.

[15] V. Cicourel (1986) The Reproduction of Objective Knowledge: Common Sense Reasoning in Medical Decision Making. *The Knowledge Society: The growing impact of scientific knowledge on social relations*. Reidel, Dordrecht.

[16] H.M. Collins & R. Evans (2002) The Third Wave of Science Studies - Studies of Expertise and Experience. *Social Studies of Science*, 32(2): 235-296.

[17] T. Davis & G. Gregory (2003) Creating Diderot Unites - quest for possible selves? *Journal of Marketing*, 20(1) :44-54.

[18] T.G. Dietterich (2000) Ensemble Methods in Machine Learning. *Int. Workshop on Multiple Classifier Systems*, LNCS, vol. 1857, Springer-Verlag, pp. 1-15.

[19] A. Dijksterhuis, M.W. Bos, L.F. Nordgren, & R.B. van Baaren (2006) Complex choices better made unconsciously? *Science*, 313, 760-761.

[20] M. Eastlik & S. Lotz (1999) Profiling potential adopters and non-adopters of an interactive electronic shopping medium. *International Journal of Retail & Distribution Management*, 27(6) :209-223.

[21] G. Elbaz (1992) *The Sociology of AIDS Activism: The Case of ACT UP/New York, 1987-1992*. PhD Dissertation, Department of Sociology, City University of New York.

[22] S. Epstein (1995) The Construction of Lay Expertise: AIDS Activism and the Forging of Credibility in the Reform of Clinical Trials. *Science, Technology & Human Values*, 20(4): 408-437.

[23] J.St.B.T. Evans (2008) Dual-processing accounts of reasoning, judgement and social cognition. *Annual Review of Psychology*, 59, 255-278.

[24] J. Fodor (1975) *The Language of Thought*. Cambridge, MA: The MIT Press.

[25] J. Fodor and E. Lepore (1996) The red herring and the pet fish: Why concepts still can't be prototypes. *Cognition*, 58, 253-270.

[26] J. Fodor (1998) *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford.

[27] H. von Foerster (1981) *Observing Systems: Selected Papers of Heinz von Foerster*. Seaside, CA: Intersystems Publications.

[28] H. von Foerster (2002) *Understanding understanding*. Springer-Verlag.

[29] E. Friedson (1988) *Profession of Medicine: A Study of the Sociology of Applied Knowledge*. University of Chicago Press.

[30] S. Funtowicz and J.R. Ravetz (1993) Science for the Post-Normal Age. *Futures*, 25(7): 735-755.

[31] P. Gärdenfors (2000) *Conceptual Spaces*. MIT Press.

[32] G. Gigerenzer (2007) *Gut Feelings - The Intelligence of the Unconscious*. Penguin Books.

[33] S. Gigerenzer, P. Todd, and the ABC Research Group (1999) *Simple heuristics that make us smart*. New York: Oxford University Press.

[34] E.M. Gold (1967) Language Identification in the Limit. *Information and Control*, 10(5): 447-474.

[35] D.G. Goldstein & G. Gigerenzer (2002) Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109, 75-90.D

[36] K. Hakkarainen, T. Palonen, S. Paavola & E. Lehtinen (2004) *Communities of networked expertise: Professional and educational perspectives*. Elsevier, Amsterdam.

[37] C.M. Harrison, J. Burgess & J. Clark J (1998) Discounted Knowledges: Farmers' and Residents' Understandings of Nature Conservation Goals and Policies. *Journal of Environmental Management*, 54: 305-320.

[38] M. van Heeswijk, Y. Miche, T. Lindh-Knuutila, T. Honkela, E. Oja & A. Lendasse (2009) Adaptive Ensemble Models of Extreme Learning Machines for Time Series Prediction. *Proc. of ICANN'09, International Conference on Artificial Neural Networks*, Lecture Notes in Computer Science, Springer, pp. 305-314.

[39] von Hippel, E., (2005) *Democratizing Innovation*. MIT Press, Cambridge (Mass.)

[40] D. Hofstadter (1995) *Fluid Concepts and Creative Analogies*. NY: Basic Books

[41] T. Honkela (1993) Neural nets that discuss: a general model of communication based on self-organizing maps. *Proc. of ICANN'93*, pp. 408-411, Springer, London.

[42] T. Honkela, V. Pulkki & T. Kohonen (1995) Contextual relations of words in Grimm tales analyzed by self-organizing map. Proc. of ICANN-95, International Conference on Artificial Neural Networks, F. Fogelman-Soulié and P. Gallinari (eds), vol. 2, EC2 et Cie, Paris, pp. 3-7.

[43] T. Honkela & J. Winter (2003) *Simulating language learning in community of agents using self-organizing maps.* Computer and Information Science Report A71, Helsinki University of Technology, Helsinki, Finland.

[44] T. Honkela, V. Könönen, T. Lindh-Knuutila & M.-S. Paukkeri (2008) Simulating processes of concept formation and communication. *Journal of Economic Methodology*, 15(3):245-259.

[45] K. Hornik, M. Stinchcombe & H. White (1989) Multilayer feedforward networks are universal approximators. *Neural Networks*, 2, 359-366.

[46] J. Hukkinen (2008) *Sustainability Networks: Cognitive Tools for Expert Collaboration in Social-Ecological Systems.* Routledge, London.

[47] D. Hume (1740) A Treatise of Human Nature. Downloaded from $http://www.gutenberg.org/dirs/etext03/trthn10.txt$, 13th of November 2009.

[48] H. Hyötyniemi (2009) Humans in Loops – a Neocybernetic View at Complex Processes. *Proc. of Automaatio XVIII 2009.*

[49] S. Hyysalo (2006) Representations of Use and Practice-Bound Imaginaries in Automating the Safety of the Elderly. *Social Studies of Science*, 36(4): 599-626.

[50] A. Irwin (1995) *Citizen Science: A Study of People, Expertise and Sustainable Development.* Routledge.

[51] A. Irwin & B. Wynne (eds.) (1996) *Misunderstanding Science? The Public Reconstruction of Science and Technology.* Cambridge University Press , Cambridge and New York.

[52] R.A. Jacobs, M.I. Jordan, S.J. Nowlan & G.E. Hinton (1991) Adaptive mixtures of local experts. *Neural computation*, 3, 79-87.

[53] N. Janasik, T. Honkela & H. Bruun (2009) Text mining in qualitative research: Application of an unsupervised learning method. *Organizational Research Methods*, 12(3): 436-460.

[54] S. Jasanoff (2003) Breaking the Waves in Science Studies: Comment on H.M. Collins an Robert Evans, The Third Wave of Science Studies. *Social Studies of Science*, 33(3): 389-400.

[55] S. Jasanoff (2003) Technologies of Humility: Citizen Participation in Governing Science. Minerva, 41: 223-244.

[56] F. Karlsson (2009) Origin and maintenance of clausal embedding complexity. G. Sampson, D. Gil and P. Trudgill (eds.), *Language Complexity as an Evolving Variable*. Oxford University Press, pp. 192-202.

[57] D. Kirsh (2003) Implicit and Explicit Representation. *Encyclopedia of Cognitive Science.*

[58] D. Kirsh (2008) Problem Solving and Situated Cognition. *Handbook of Situated Cognition*. Cambridge University Press, pp. 264-306.

[59] T. Kohonen (2001) *Self-Organizing Maps*. Springer.

[60] T. Kotro (2005) *Hobbyist knowing in product development: desirable objects and passion for sports in Suunto Corporation*. Helsinki: University of Art and Design Helsinki.

[61] R. Kowalski and D. Kuehner (1971) Linear Resolution with Selection Function. *Artificial Intelligence*, 2(3/4): 227-260.

[62] K. Lagus, A. Airola and M. Creutz (2002) Data analysis of conceptual similarities of Finnish verbs. *Proc. of CogSci 2002*, pp. 566-571.

[63] R. Langlais, N. Janasik & H. Bruun (2004) Managing Knowledge Network Processes in the Commercialization of Science: Two Probiotica Discovery Processes in Finland and Sweden. *Science Studies*, 17(1): 34-56.

[64] S. Laurence & E. Margolis (2002) Radical concept nativism. *Cognition*, 86(1): 25-55.

[65] E. Lazega (2001) *The collegial phenomenon: The social mechanisms of cooperation among peers in a corporate law partnership*. England: Oxford University Press.

[66] T. Lindh-Knuutila, T. Honkela & K. Lagus (2006) Simulating meaning negotiation using observational language games. *Symbol Grounding and Beyond*, Springer, pp. 168-179.

[67] T. Lindh-Knuutila, J. Raitio, and T. Honkela (2009) Combining self-organized and Bayesian models of concept formation. *In Proc. of NCPW11.*

[68] R.P. Lippmann (1989) Pattern classification using neural networks. *IEEE Communications magazine*, 27(11): 47-54.

[69] B.S.-C. Liu, R. Madhavan & D. Sudharshan (2005) DiffuNET: The impact of network structure on diffusion of innovation. *European Journal of Innovation Management*, 8(2): 240-262.

[70] N. Luhmann (1995) Social Systems, Stanford: Stanford University Press.
(Originally: Soziale Systeme: Grundriss einer allgemeinen Theorie, Frankfurt: Suhrkamp, 1984)

[71] D. Merritt (1989) *Building Expert Systems in Prolog*. Springer-Verlag.

[72] C.A. Mitchell, A.L. Carew & R. Clift (2004) The role of the professional engineer and scientist in sustainable development. *Sustainable Development in Practice: Case Studies for Engineers and Scientists*. John Wiley and Sons, Chichester.

[73] J. Nahapiet & S. Ghostal (1998) Social capital, intellectual capital, and the organizational advantage. *Academy of Management Review*, 22(2): 242-266.

[74] I. Nishizaki, H. Katagiri & T. Oyama (2009) Simulation Analysis Using Multi-Agent Systems for Social Norms. *Computational Economics*, 34(1): 37-65.

[75] I. Nonaka & H. Takeuchi (1995) *The knowledge creating company: how Japanese companies create the dynamics of innovation*. New York: Oxford University Press.

[76] E. Oja (2002) Unsupervised learning in neural computation. *Theoretical Computer Science*, 287: 187-207.

[77] M. Pantzar & E. Shove (2009) Understanding innovation in practice: a discussion of the production and re-production of Nordic Walking. *Technology Analysis and Strategic Management*, Volume 22, issue 2, forthcoming.

[78] M. Polanyi (1966) *The Tacit Dimension*. Doubleday & Co.

[79] R. Polikar (2009) Ensemble learning. *Scholarpedia*, 4(1): 2776.

[80] I. Pyysiäinen, M. Lindeman & T. Honkela (2003) Counterintuitiveness as the hallmark of religiosity. *Religion*, 33(4): 341-55.

[81] J. Raitio, R. Vigário, J. Särelä & T. Honkela (2004) Assessing similarity of emergent representations based on unsupervised learning. *Proc. of IJCNN 2004*, vol. 1., pp. 597-602.

[82] J.R. Ravetz (1973) *Scientific Knowledge and its Social Problems*. Penguin Books, Middlesex.

[83] O. Renn (2001) The need for integration: risk policies require the input from experts, stakeholder and the public at large. *Reliability engineering & System Safety*, 72: 131-135.

[84] O. Renn, T. Webler, H. Rakel, P. Dienel & B. Johnson (1993) Public-Participation in Decision-Making – A 3-Step Procedure. *Policy Sciences*, 26: 189-214.

[85] P. Repo, M. Pantzar, P. Timonen & K. Hyvönen (2006) Blueprint for Consumer 2.0. P. Walden, R. Fullér & J. Carlsson (Eds.): Expanding the Limits of the Possible, pp. 94-100. IAMSR: Turku.

[86] H. Ritter & T. Kohonen (1989) Self-organizing semantic maps. *Biological cybernetics*, 61: 241-254.

[87] C. Roberts (2006) "What Can I Do to Help Myself?" Somatic Individuality and Contemporary Hormonal Bodies. *Science Studies*, 19(2): 54-76.

[88] E. Rogers (1995) *Diffusion of innovations*, fourth ed. Free Press, New York.

[89] A.-M. Rusanen, O. Lappi, T. Honkela & M. Nederström (2008) Conceptual coherence in philosophy education - visualizing initial conceptions of philosophy students with self-organizing maps. *Proc. of CogSci'08*, pp. 64-70.

[90] J. Salojärvi, K. Puolamäki, E. Savia & S. Kaski (2008) *Inference with Discriminative Posterior*. Manuscript, downloaded from http://arxiv.org/pdf/0807.3470

[91] J. Santrock (2004) Cognitive Development Approaches (Chapter 6). *A Topical Approach To Life-Span Development*. McGraw-Hill, New York, NY, pp. 200-225.

[92] G. Schwenk & T. Reimer (2008) Simple Heuristics in Complex Networks: Models of Social Influence. Journal of Artificial Societies and Social Simulation, 11(3):4 ($http : //jasss.soc.surrey.ac.uk/11/3/4.html$)

[93] Y. Shoham & K. Leyton-Brown (2009) *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. New York: Cambridge University Press.

[94] E. Shove & M. Pantzar (2007) Recruitment and Reproduction: The careers and carriers of digital photography and floorball. *Human Affairs*, 17, 164-167.

[95] E. Shove, M. Watson, M. Hand & J. Ingram (2007) *The Design of Everyday Life*. Berg.

[96] H.T. Siegelmann (1999) *Neural Networks and Analog Computation: Beyond The Turing Limit*. Boston, MA: Birkhauser.

[97] H.A. Simon (2002) Near decomposability and the speed of evolution. *Industrial and Corporate Change*, 11(3): 587-599.A

[98] Z. Solan, D. Horn, E. Ruppin, and S. Edelman, S. (2005) Unsupervised learning of natural languages. *Proceedings of the National Academy of Sciences*, 102(33): 11629-11634

[99] E. Soloway, B. Adelson & K. Ehrlich (1988) Knowledge and Processes in the Comprehension of Computer Programs. M.T.H. Chi, R. Glaser and M.J. Farr (eds.), *The Nature of Expertise*, pp. 129-152, Hillsdale, New Jersey: Lawrence Erlbaum Associates.

[100] S.L. Star & J.R. Griesemer (1989) Institutional Ecology, 'Translations,' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907 - 1939. *Social Studies of Science*, 19: 387-420.

[101] L. Steels (1998) The Origins of Ontologies and Communication Conventions in Multi-Agent Systems. *Autonomous Agents and Multi-Agent Systems*, 1(2): 169-194.

[102] L. Sterling & E. Shapiro (1994) *The Art of Prolog*. (2nd edition). The MIT Press.

[103] R. Sun (2006) *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. Cambridge University Press.

[104] J. Venna, S. Kaski, & J. Peltonen (2003) Visualizations for assessing convergence and mixing of MCMC. In N. Lavrac, D. Gamberger, H. Blockeel, and L. Todorovski (eds.), *Proc. of the 14th European Conference on Machine Learning (ECML 2003)*, pp. 432-443, Springer, Berlin.

[105] P. Vogt (2005) The emergence of compositional structures in perceptually grounded language games. *Artificial Intelligence*, 167(1-2): 206-242.

[106] W. Weaver (1948) Science and Complexity. *American Scientist*, 36: 536.
(Available at $http://www.ceptualinstitute.com/genre/weaver/weaver-1947b.htm$)

[107] L. Wittgenstein (1953) *Philosophical Investigations*. Blackwell, Oxford.

[108] D. Wolpert (1996) The Lack of A Priori Distinctions between Learning Algorithms. *Neural Computation*, 8(7): 1341-1390.

[109] D.H. Wolpert & W.G. Macready (1997) No Free Lunch Theorems for Optimization. *IEEE Transactions on Evolutionary Computation*, 1(1): 67-82.

[110] B. Wynne (1996) May the Sheep Safely Graze? A Reflexive View of the Expert-Lay Knowledge Divide. *Risk, Environment and Modernity: Towards a New Ecology*. Sage, London.

[111] J. Zhang (2009) Adaptive learning via selectionism and Bayesianism, Part I: Connection between the two. *Neural Networks*, 22(1): 220-228.

TKK REPORTS IN INFORMATION AND COMPUTER SCIENCE

TKK-ICS-R14	Jori Dubrovin Checking Bounded Reachability in Asynchronous Systems by Symbolic Event Tracing. April 2009.
TKK-ICS-R15	Eerika Savia, Kai Puolamäki, Samuel Kaski On Two-Way Grouping by One-Way Topic Models. May 2009.
TKK-ICS-R16	Antti E. J. Hyvärinen Approaches to Grid-Based SAT Solving. June 2009.
TKK-ICS-R17	Tuomas Launiainen Model checking PSL safety properties. August 2009.
TKK-ICS-R18	Roland Kindermann Testing a Java Card applet using the LIME Interface Test Bench: A case study. September 2009.
TKK-ICS-R19	Kalle J. Palomäki, Ulpu Remes, Mikko Kurimo (Eds.) Studies on Noise Robust Automatic Speech Recognition. September 2009.
TKK-ICS-R20	Kristian Nybo, Juuso Parkkinen, Samuel Kaski Graph Visualization With Latent Variable Models. September 2009.
TKK-ICS-R21	Sami Hanhijärvi, Kai Puolamäki, Gemma C. Garriga Multiple Hypothesis Testing in Pattern Discovery. November 2009.
TKK-ICS-R22	Antti E. J. Hyvärinen, Tommi Junttila, Ilkka Niemelä Partitioning Search Spaces of a Randomized Search. November 2009.
TKK-ICS-R23	Matti Pöllä, Timo Honkela, Teuvo Kohonen Bibliography of Self-Organizing Map (SOM) Papers: 2002-2005 Addendum. December 2009.

ISBN 978-951-22-262-4 (Print)
ISBN 978-951-22-263-1 (Online)
ISSN 1797-5034 (Print)
ISSN 1797-5042 (Online)