Knowledge Representation and the Embodied Mind: Towards a Philosophy and Technology of Personalized Informatics

Wolf-Tilo Balke[1], Klaus Mainzer[2]

[1] Electrical Engineering & Computer Science,
University of California, Berkeley, CA 94720, USA
balke@eecs.berkeley.edu, http://www.cs.berkeley.edu/~balke
[2] Chair for Philosophy of Science, Institute of Interdisciplinary Informatics,
Universität Augsburg, 86135 Augsburg, Germany
klaus.mainzer@phil.uni-augsburg.de,
http://www.informatik.uni-augsburg.de/I3/

Abstract. Knowledge representation has a long tradition in logic and philosophy. Automated reasoning with ontologies and categories had been discussed in philosophy, before it was formalized in artificial intelligence and e.g. applied in information systems. But, most of our knowledge is implicit and unconscious, situated and personalized. It is not formally represented, but embodied knowledge, which is learnt by doing, applied by self-organization, and understood by bodily interacting with social environments. In a complex world, we must act and decide with incomplete and fuzzy knowledge under the conditions of bounded rationality. The bounded rationality of embodied minds is a challenge of informatics especially in the complex information world of Internet applications and Web services. It overcomes traditional concepts of mind-body dualism in the philosophy of mind, traditional knowledge representation in AI, and rational agents („homo oeconomicus“) in economics. Personalized informatics opens a trans-disciplinary perspective for philosophy and working technology.

1 Knowledge and Representation

Knowledge representation which is today used in database applications, artificial intelligence, software engineering, and many other disciplines of computer science has deep roots in logic and philosophy [15]. In the beginning, there was Aristotle (384-322 B.C.) who developed logic as a precise method for reasoning about knowledge. Syllogisms were introduced as formal patterns for representing special figures of logical deductions. According to Aristotle, the subject of ontology is the study of categories of things that exist or may exist in some domain. Aristotle distinguished ten basic categories for classifying anything that may be said or predicated about anything: Substance, quality, quantity, relation, activity, passivity, having, situatedness, spatiality, and temporality. Many of these categories are today applied in, e.g. databases. In the middle ages, knowledge representation was illustrated by graphic diagrams and pictures. In the ‘summulae logicales’ (1239) of Peter of Spain, an ontological hierarchy with aristotelian categories represented knowledge by genus (supertype) and species (subtype). The features that distinguished different species of the same genus were called differentiae. Raimundus Lullus (13th century) illustrated an ontological hierarchy by a tree with branches for categories. Leaves corresponded to questions or to answers which should automatically be found by a system of rotating disks for combining features of things. Actually, Raimundus Lullus applied a kind of British Museum algorithm, the first attempt to develop mechanical aids for problem solving and information retrieval. Today, we use Entity-Relationship (ER)-diagrams in suitable forms to illustrate structures of ontologies in informatics.

In modern times, Descartes considered the human brain as a store of knowledge representation. Recognition was made possible by an isomorphic correspondance between internal geometrical representations (ideae) and external situations and events. Leibniz was deeply influenced by these traditions. In his 'mathesis universalis', he required a universal formal language (lingua universalis) to represent human thinking by calculation procedures and to implement them to mechanical calculating machines. An 'ars iudicandi' should allow every problem to be decided by an algorithm after representation in numeric symbols. An 'ars iveniendi' should enable users to seek and enumerate desired data and solutions of problems. Thus, in the age of mechanics, knowledge representation was reduced to mere mechanical calculation procedures. In Kant's epistemology, recognition is not only a passive mapping of the external world, but an active construction of internal representations by a priori categories of pure reason. In modern terms, categories are considered as tools which must be assumed before ('a priori') any application of knowledge representation. Cognitive constructivism roots back to Kant's epistemology. In the tradition of Brentano's and Husserl's phenomenology, Aristotelian ontologies had been discussed for knowledge representation, again. Recognition needs intentional actions, which direct our awareness and consciousness to objects of the world. Thus, according to Husserl, understanding is not possible by symbolic representations of the external world, but by the intentionality of human consciousness [7]. Intentionality became a prominent criterion to distinguish human consciousness and computer representation of knowledge in recent AI-debates [18].

Computational cognitivism arose on the background of Turing's theory of computability. In his functionalism, the hardware of a computer is related to the wetware of the human brain. The mind is understood as the software of a computer. Turing argues: If human mind is computable, it can be represented by a Turing program (Church's thesis) which can be computed by a universal Turing machine, i.e. technically by a general purpose computer. Even if people do not believe in Turing's strong AI-thesis, they often claim classical computational cognitivism in the following sense: Computational processes operate on symbolic representations referring to situations in the outside world. These formal representations should obey Tarski's correspondence theory of truth: Imagine a real world situation X1 (e.g., some boxes on a table) which is encoded by a symbolic representation A1 = encode(X1) (e.g., a description of the boxes on the table). If the symbolic representation A1 is decoded, then we get the real world situation X1 as its meaning, i.e. decode(A1) = X1. A real-world operation T (e.g. a manipulation of the boxes on the table by hand) should produce the same real-world result A2, whether performed in the real world or on the symbolic representation: decode(encode(T)(encode(X1))) = T(X1) = X2. Thus, there is an isomorphism between the outside situation and its formal representation in Cartesian tradition.

2 Self-Organization and the Embodied Mind

Knowledge representations with ontologies, categories, frames, and scripts of expert systems work along the discussion in section 1. However, they are restricted to a specialized knowledge base without the background knowledge of a human expert. Human experts do not rely on explicit (declarative) rule-based representations, but on intuition and implicit (procedural) knowledge [6]. Further on, as already Wittgenstein knew, our understanding depends on situations. The situatedness of representations is a severe problem of informatics. A robot, e.g., needs a complete symbolic representation of a situation which must be updated if the robot's position is changed. Imagine that it surrounds a table with a ball and a cup on it. A formal representation of their respective relative positions in a computer language may be ON(TABLE, BALL), ON(TABLE, CUP), BEHIND(CUP, BALL), etc. Depending on the robot's position relative to the arrangement, the cup is sometimes behind the ball or not. So, the formal representation BEHIND(CUP, BALL) always has to be updated in changing positions. How can the robot prevent incomplete knowledge? How can it distinguish between reality and its relative perspective? Situated agents like human beings need no symbolic representations and updating. They look, talk, and interact bodily, e.g., by pointing to things. Even rational acting in sudden situations does not depend on internal representations and logical inferences, but on bodily interactions with a situation (e.g. looking, feeling, and reacting).

Thus, we distinguish formal and embodied acting in games with more or less similarity to real life: Chess for instance is a formal game with complete representations, precisely defined states, board positions, and formal operations. Soccer is a non-formal game with skills depending on bodily interactions, without complete representations of situations and operations, which are never exactly identical. According to Merleau-Ponty, intentional human skills do not need any internal representation, but they are trained, learnt, and embodied in an optimal 'gestalt', which cannot be repeated [16]. An athlete like a pole-vaulter cannot repeat his/her successful jump like a machine generating the same product. Husserl's representational intentionality is replaced by embodied intentionality. The embodied mind is no mystery. Modern biology, neural, and cognitive science give many insights into its origin during the evolution of life. The key-concept is self-organization of complex dynamical systems [13]. Brains are neural systems which allow quick adaptation to changing situations during the life-time of an organism. In short: They can learn, assess and anticipate. The human brain is a complex system of neurons self-organizing in macroscopic patterns by neurochemical interactions. Perceptions, emotions, thoughts, and consciousness correspond to these neural patterns. Motor knowledge for instance is learnt in an unknown environment and stored implicitly in the distribution of synaptic weights of the neural nets. In the human organism, e.g. walking is a complex bodily self-organization, largely without central control of brain and consciousness: It is driven by the dynamical pattern of a steady periodic motion, the attractor of the motor system. Motor intelligence emerges without internal symbolic representations.

But not only ‘low level’ motor intelligence, but also ‘high level’ cognition (e.g., categorization) can emerge from complex bodily interaction with an environment by sensory-motor coordination without internal symbolic representation. We call it ‘embodied cognition’: An infant learns to categorize objects and to build up concepts by touching, grasping, manipulating, feeling, tasting, hearing, and looking at things, and not by explicit representations. The categories are based on fuzzy patchworks of prototypes and may be improved and changed during life. We have an innate disposition to construct and apply conceptual schemes and tools (in the sense of Kant). Moreover, cognitive states of persons depend on emotions. We recognize emotional expressions of human faces with pattern recognition of neural networks and react by generating appropriate facial expressions for non-verbal communication. Emotional states are generated in the limbic system of the brain which is connected with all sensory and motoric systems of the organism. All intentional actions start with an unconscious impulse in the limbic system which can be measured half a second before their performance. Thus, embodied intentionality is a measurable feature of the brain [8]. Humans use feelings to help them navigate the ontological trees of their concepts and preferences, to make decisions in the face of increasing combinatorial complexity: Emotions help to reduce complexity.

The embodied mind is obviously a complex dynamical system acting and reacting in dynamically changing situations. The emergence of cognitive and emotional states is made possible by brain dynamics which can be modeled by neural networks. According to the principle of computational equivalence [13, 14], any dynamical system can be simulated by an appropriate computational system. But, contrary to Turing’s AI-thesis, that does not mean computability in any case. In complex dynamical systems, the rules of locally interacting elements (e.g., Hebb’s rules of synaptic interaction) may be simple and can be programmed in a computer model. But their nonlinear dynamics can generate complex patterns and system states, which cannot be predicted in the long run without increasing loss of computability and information. Thus, artificial minds could have their own intentionality, cognitive and emotional states that cannot be forecast and computed like in the case of natural minds [5]. Limitations of computability are characteristic features of complex systems.

In a complex dynamical world, decision-making and acting is only possible under conditions of bounded rationality. Bounded rationality results from limitations on our knowledge, cognitive capabilities, and time. Our perceptions are selective, our knowledge of the real world is incomplete, our mental models are simplified, and our powers of deduction and inference are weak and fallible. Emotional and subconscious factors affect our behavior. Deliberation takes time and we must often make decisions before we are ready. Thus, knowledge representation must not be restricted to explicit declarations. Tacit background knowledge, change of emotional states, personal attitudes, and situations with increasing complexity are challenges of modeling information and communication systems. Personalized information systems in dynamic situations should be referred to ubiquitous and invisible computing of world-wide interactive media, in order to improve human-oriented information services and to support a sustainable information world.

3 Towards Advanced Personalization in Computer Systems

Especially for areas in computer science that rely on a user's expression of individual needs like e.g. query processing in databases and information systems, media retrieval in document collections or selection problems in Web services or e-commerce workflows, getting a precise account of each user's profile/preferences is mission critical. But since research in psychology shows that even in purposeful tasks users are usually not fully conscious of their exact wishes and needs, see e.g. [1], eliciting preferences directly from users is a difficult matter. It often needs a tedious process like the manual selection of services or areas of interest for personalization in publish/subscribe systems. Moreover, given that some knowledge is embodied the elicited information will be naturally incomplete and simply logging and storing and using user-stated keywords/behaviour will sometimes lead to counterintuitive results. In order to raise a personalized system's performance in terms of relevance, a system thus has not only to focus on explicit user specification, but should also take information into account, that is specified by the user's implicit notions, situation or assumed common knowledge. This information can be gathered mainly from four sources:

- **long-term preferences:** The notion of relevance from previous interactions or generally applicable knowledge about a user is used
- **intention:** The specific user's purpose of the interaction is included in personalizing the system
- **situation:** The present state and environment of a user is used to decide whether specific preferences or rules are applicable
- **domain:** Knowledge on the specific domain (often referred to as expert knowledge) is used within an interaction

Let us consider typical instances of these kinds of personalization information. Among long-term preferences typical re-occurring individual preferences are collected e.g. individual tastes like colours, general areas of interest or preferred layout settings. Generally this kind of preference can always be used to personalize a system for individual users and is the usual kind often stored in user profiles. Systems, however, cannot always rely on these preferences, since they might be either further specified for certain categories or simply not applicable in a certain context. Consider for instance a set of colour preferences in an e-commerce setting. Though a user can be assumed to have a certain favourite colour that will apply to shopping decisions, the preferences might be different for e.g. clothing and cars, since driving an e.g. red car differs from actually wearing red clothing. Moreover, for e.g. book shopping the colour preference becomes entirely inapplicable, since the request to by a red is usually not sensible. A basic framework for eliciting and tagging this kind of information can be found in e.g. [10].

Of a less general, but more interesting kind for personalization tasks are the preferences for the last three categories. The (assumed) intention of a user will help to decide, which choices a user should be offered in personalization and a user generally cares about at a certain point in time during his/her interaction with the system. Typical examples for using these preferences are e.g. adaptive hypertext applications, where depending on a user's previous interactions or navigation patterns the environment can be personalized, see e.g. [11]. Also the situation has an impact on how to personalize a system. Context-aware systems use clues from a user's direct environment (like time or location), personal characteristics like emotional states, technical characteristics like client device capabilities, or certain high level information like "user in a business meeting" or "user at home" for personalization tasks. Examples for systems integrating this kind of information are location-based services or situation-based communication routing, or context-aware synthesis of multimedia content like discussed in e.g. [19]. The most renowned realizations of the last kind of preferences for personalization are the so-called expert systems that encode domain knowledge elicited from domain expert in a system. Since [6] shows that there cannot be a complete set of expert knowledge rules, since most expert knowledge is not represented by rules, but embodied in the expert, we cannot consider domain preferences in the sense of expert systems, but have to rely on domain specific heuristics like which general preferences (and in what combination) might be applicable, or what users generally care about in a certain domain.

In today's systems the latter three kinds of preferences – if at all – are mostly built directly into the application logic and represent the embodied mind as opposed to collected individual long-term preferences. In the next section we will consider two sample scenarios and focus on how to use these kinds of preferences not in a hard-coded fashion, but flexibly mixed with information from user profiles.

4 Case Studies for Preference-based Personalization

Let us consider two short case studies where we can see parts of the embodied mind represented in preferences. For effective personalization knowledge from all four sources discussed above has to be blended with the specific user-provided details/keywords for an interaction. Though generally not all embodied knowledge can be captured that way, this method nevertheless provides a useful way of personalizing systems under the notion of bounded rationality.

First consider personalized retrieval tasks in databases and information systems. As stated in [2], expanding queries along user-specific preferences goes back to the area of cooperative answering, see e.g. [17]. The basic notion of cooperative retrieval systems is that they will relax the user-specified terms until a match in a collection of data can be found. Thus even an overspecified query will lead to some 'best efforts' results and avoid empty result sets and the often necessary refinement of queries. This way of dealing with query predicates as soft constraints, is also necessary for personalization tasks using individual preferences that have not been explicitly stated for a specific interaction. Since they have been implicitly assumed by the system as representations of common or embodied knowledge either from long-term profiles, intention, situation, or domain, they have to be considered on a lower level (i.e. as soft constraints that may refine too large result sets, but are relaxed if empty result sets are retrieved). Recently [12] introduced a system of integrating preferences in the form of strict partial order with an "I like A better than B" semantics into database queries. Here basic preferences can be modelled and combined into more complex queries using operators for deriving Pareto sets (i.e. all preferences are considered as equally important), prioritized sets (i.e. a certain order is imposed on the preferences) and ranked result sets (i.e. preferences on numerical domains are aggregated using utility functions).

As a second example let us consider discovering useful Web services or selecting suitable services to construct complex workflow requests in a personalized manner like shown in [3, 4]. When designing a service like restaurant reservation or flight booking service providers usually already have quite specific ideas what capabilities the service should provide and what kinds of interaction to expect. Thus providers are domain experts, who can provide a set of useful domain preferences and even ontologies (that can be understood as categorizations encoding common knowledge) to foster successful execution/composition of services even for non-expert users. Moreover, providers also may anticipate different possibilities for usage of the service (possibly also in different situation scenarios). Generally in well-defined services only a certain number of typical requests/business processes will exist. These typical interactions for different users/groups also are preference patterns or usage patterns open for our personalization approach. A usage pattern may e.g. depend on the basic intentions of a significant group of users. Different intentions will need different patterns that reflect on both a user's profile stating his or her notion of a service's usefulness or desired characteristics (like execution costs, quality guarantees, etc.) and the service profiles that are employed to carry out the actual business task. Also here the basic method of relaxing demands is necessary to answer user requests in a cooperative fashion.

5 Summary and Outlook

In this paper we focused on the necessary representation of knowledge for personalization tasks in Informatics. Starting from the notion that most relevant information for personalization tasks cannot entirely be elicited as expert knowledge, but is mostly embodied in the individual user (which is also consistent with current brain research), we propose to use flexible preference-based frameworks to personalize computer systems under the paradigm of bounded rationality.

As shown for typical user interaction in the areas of personalized retrieval in databases/information systems and proactive Web service discovery/selection personalizing the interaction with preferences from each individual user's long-term profile, intention, situation, and domain promises to result in an improved information exchange with (and thus improved usability of) computer systems and allows for better service provisioning. This is because using individual profiles the user-provided information can be expanded with information representing the 'embodied' information necessary for a certain task. Since this information is not conscious, this expansion really adds value to the personalized task. However, since all preferences used for expanding the user information are only used on a lower level of importance than the explicitly provided information (and will be relaxed if necessary), our expansion will respect an individual users needs and never violate explicit constraints.

Obviously, personalized computer systems do not aim at complete computational models of the human embodied mind, which was an impractical illusion of traditional

AI and expert systems research. In the current trend of modern informatics, we want to construct effective and appropriate tools and service systems under the conditions of bounded knowledge/rationality which need interdisciplinary cooperation between informatics, cognitive science and philosophy of mind.

References

1. Ackhoff, R., Emery, F.: On Purposeful Systems. Aldine Atherton, Chicago (1972)
2. Balke, W.-T.: A Roadmap to Personalized Information Systems by Cognitive Expansion of Queries. Int. Workshop on Content-Knowledge Management and Mass Personalization of E-Services (EnCKompass 2002), Paris, France (2002)
3. Balke, W.-T., Wagner, M.: Cooperative Discovery for User-centered Web Service Provisioning. Int. Conf. on Web Services (ICWS'03), Las Vegas, USA (2003)
4. Balke, W.-T., Wagner, M.: Towards Personalized Selection of Web Services. Int. World Wide Web Conference (WWW'03), Budapest, Hungary (2003)
5. Dennett, C.D.: Brainchildren: Essays on Designing Minds, MIT Press1998
6. Dreyfus, H.L.: What Computer's can't do – The Limits of Artificial Intelligence. Harper & Row, New York, USA (1979)
7. Dreyfus, H.L.: Husserl, Intentionality, and Cognitive Science. MIT Press, Cambridge, USA (1982)
8. Freeman, W.J.: How and why brains create meaning from sensory information. Int. J. Bifurcation and Chaos 14 (2004) 515-530
9. Floridi, L. (ed.): Philosophy of Computing and Information. Blackwell, Oxford, UK (2004)
10. Holland, S., Kießling, W.: Situated Preferences and Preference Repositories for Personalized Database Applications. Int. Conf. on Conceptual Modeling (ER'04), Shanghai, China (2004)
11. Kaplan, C., Fenwick, J., Chen, J.: Adaptive Hypertext Navigation Based On User Goals and Context. Int. J. on User Modeling and User-Adapted Interaction. Vol. 3(3) (1993)
12. Kießling, W.: Foundations of Preferences in Database Systems. Int. Conf. on Very Large Data Bases (VLDB'02). Hong Kong, China (2002)
13. Mainzer, K.: Thinking in Complexity. The Computational Dynamics of Matter, Mind, and Mankind. 4th enlarged ed. Springer, New York, USA (2004)
14. Mainzer, K.: KI - Künstliche Intelligenz. Grundlagen intelligenter Systeme. Wissenschaftliche Buchgesellschaft, Darmstadt, Germany (2003)
15. Mainzer, K.: Computerphilosophie – Zur Einführung. Junius Verlag, Hamburg, Germany (2003)
16. Merleau-Ponty, M.: Phenomenology of Perception. Routledge & Kegan Paul (1962)
17. Minker, J.: An Overview of Cooperative Answering in Databases. Int. Conf. on Flexible Query Answering Systems (FQAS'98), Springer LNCS 1495, Roskilde, Denmark (1998)
18. Searle, J.R.: Intentionality. An Essay in the Philosophy of Mind. Cambridge University Press (1983)
19. Wagner, M., Balke, W.-T., Kießling, W.: An XML-based Multimedia Middleware for Mobile Online Auctions. In *Enterprise Information Systems III*, Kluwer, The Netherlands (2002)