Integrated Reasoning Systems
- A Knowledge-Level Perspective with a Case Based Bias

Agnar Aamodt
Norwegian University of Science and Technology,
Dept. of Computer and Information Science
N-7491, Trondheim, Norway

Abstract
An increasing amount of work in AI in general, and case-based reasoning in particular, is addressing means of integrating different types of reasoning methods for building and maintaining AI systems. This paper first takes a step back in an attempt to identify a descriptive and comparative framework for the corresponding methods. It is argued that the knowledge level is the appropriate level for describing the behaviour of an intended system, and for identifying its methods and knowledge components. Main research activities in the author's research group in Trondheim over the last twenty years are summarized with reference to the framework.

1. Introduction

AI is becoming a mature field of computer science. It also has a rather long history compared to many other computing areas. In the AI community we have had our ups and downs, as in most young areas striving to grow up, and we have learned our lessons. Important progress has been made in understanding the capabilities and limits of the different main approaches and specific methods that constitute our field. Still, of course, we have a long way to go before our vision of artificially intelligent systems are met.

As individual methods become better understood, one of the big challenges is how to combine individual methods to achieve system behaviour beyond the reach of each single method. This is reflected in the observation that AI researchers these days to an increasing degree are addressing method combination, or integration, issues. In the Knowledge-Based Systems group of the Department of Computer and Information Science at NTNU, we have over many years contributed to this trend. Our work has been – and are – based on the following research scope and set of premises (also somewhat reflected in the long and winding title of this paper):

i) To realize intelligent computational systems, the human mind is one important model. Cognitive Science is therefore an interdisciplinary area from which many of our methods have been influenced, and to which our research should be able to contribute. In our research group are addressing application domains that are open and do not have strong domain theories, i.e. domains in which humans perform reasonably well despite the incompleteness and uncertainties involved.

ii) The symbol-processing paradigm represents methods that are necessary for achieving the type of intelligent behaviour we seek in future AI systems. Good old-fashioned AI (GOFAI) is still very much alive and continuously getting modernized. However, symbol-processing methods do not cover all needs for future AI systems. So-called "sub-symbolic" or "non-symbolic" AI represents useful additional method areas for the learning and realization of some type of behaviour.

iii) An important feature of human reasoning is our ability to capture and remember specific situations that we experience, often referred to as cases, and our subsequent recall and reuse of these cases in understanding new situations and solving new problems. Past experiences constitute a necessary part of any intelligent system's source

This paper was presented at NAIS 2009; see http://http://events.idi.ntnu.no/nais2009/

of knowledge, and retaining them as explicit cases enables their flexible reuse when needed. However, generalized knowledge, for example in the form of heuristic rules or deeper models, is also an important type of knowledge learned and used by humans. Our hypothesis is that the integration of both types of knowledge and associated reasoning methods in an artificially intelligent system can improve the power of each individual method. This also applies to the combination of symbol-processing and other method types.

iv) Given that integration of methods is a driving force in our research, a high level descriptive and analytic framework is needed. The Knowledge Level, as described by Allen Newell (Newell, 1982), is the appropriate level of such a framework. The development of knowledge modeling and systems description methodologies within the knowledge acquisition community has operationalized Newell's notion of the knowledge level for the purpose of analyzing and designing intelligent systems.

In the rest of this paper a knowledge-level modeling account of the type of systems we are targeting is first presented. A "model construction approach" to the reasoning process, including both problem solving and learning, is then discussed. This is followed by a summary of some of the past and present activities in our research group at NTNU. A discussion of future challenges closes the paper

2. The knowledge level as modeling framework

In Newell's paper (Newell, 1982) the *knowledge level* was proposed as a distinct level of description of computer systems, defined to lie above the level of data structures and programming languages. The latter was referred to as the *symbol level*. In Newell's framework, each computer system level has a medium of expression, identifying what is being processed at that level. Each level further has a behavioural law, which determines how the processing is done, and which enables explanation and prediction of system behaviour at that level. At the symbol level, the medium is symbols (data structures and programs), and the behavioural law is sequential interpretation of program procedures. The knowledge level has knowledge, in terms of goals and means to obtain them, as its medium, and what is called the ”principle of rationality”, as its behavioural law. A system is described at the knowledge level as an intelligent agent with its own goals and with knowledge of how to achieve its goals. The principle of rationality states that an agent always will use its knowledge in a way that ensures the achievement of its goals - provided the agent has the knowledge needed.

The knowledge level enables a system (existing or anticipated) to be described in terms of what it does and why it wants to do it, completely independent of implementational constraints. Hence it is a level applicable to any agent to which it makes sense to ascribe knowledge and rationality (e.g. humans, some animals, some computer systems, ...). The problem with using the knowledge level in this sense, for the modeling and design of computer systems, is that it has no a priori structure. Hence, the knowledge level in this sense cannot be used directly to analyze and structure knowledge. Further, the principle of rationality assumes an ideal rational agent, not bounded by physical or temporal constraints. The knowledge level in its ”pure” form, as introduced by Newell, is therefore not particularly useful for structuring a knowledge modeling effort. This has lead to modifications of the original knowledge level notion defined by Newell, into a more operational notion of the knowledge level. This may be viewed as moving the knowledge level slightly in the direction of the symbol level. Terms that have been used to characterize this intermediate level include the "knowledge use level" (Steels, 1990), and "knowledge level architecture" (Sticklen, 1989). It also includes introducing the notion of "tractable rationality" (Van de Velde, 1993), as a means to deal with the pragmatics of real world situations as opposed to Newell’s ideal, unbounded rationality.

Research within the knowledge acquisition community has produced several methodologies and techniques for modeling systems at this level of knowledge-level architectures. Influential examples are the CommonKADS methodology (Breuker and Van de Velde, 1994), the Components of Expertise framework (Steels, 1990), the Generic Tasks approach (Chandrasekaran, 1992), Role Limiting Methods (McDermott, 1988), and the Method-to-Task approach underlying the PROTEGE systems (Musen, 1989). Work in order to unify several of these methodologies has been a focus of several groups, as exemplified by the multiple perspective approach of the KREST methodology (Steels, 1993), and by the generality strived for in CommonKADS (Wielinga et. al., 1992). All these approaches have a common feature, they view knowledge modeling - at least partly - from what is claimed to be a knowledge-level perspective.

Given this refined and operationalized notion of the knowledge level, a consensus seems to have been established that knowledge should be grouped into three main types, or viewed from three perspectives: *Task knowledge*, *Method knowledge*, and *Domain knowledge* (see Figure 1 Knowledge perspectives). Task knowledge models what to do, usually in a task-subtask hierarchy. Tasks are tightly connected to goals, and the terms are sometimes used interchangeably. A task is defined by the goals that a system tries to achieve. Method knowledge describes how to do it, i.e. a method is a means to accomplish a task (e.g. to solve a problem). Domain knowledge is the knowledge about the world that a method needs to accomplish its task. Examples are facts, heuristics, causal relationships, multi-relational models, and – of course – specific cases. The term “domain knowledge” is not a particularly good term, however, since task- and method knowledge often is domain specific as well. It is hard to find a better term to indicate this type of knowledge, however, although 'object knowledge' 'application knowledge' and just 'models' have been proposed. We will stick to the term domain knowledge, but bear in mind that the other knowledge types are not necessarily domain independent. Although the in-principle decomposition at the top level into these three knowledge types is agreed upon, the naming of them, and their subdivision and interrelationships are to a large degree what characterizes each specific knowledge-level modeling methodology.

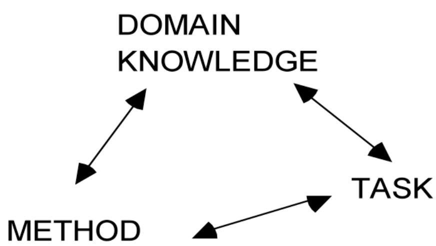

Figure 1 Knowledge perspectives

3. The model construction view of reasoning.

The model construction view of problem solving and learning comes out of a knowledge-level analysis that views reasoning itself as the construction of a model. It states that reasoning is the process of moving from an initial model instance that describes a problem to be solved, pushing this model through a set of problem solving states, finally ending with a version of the model instance that also contains a solution to the problem (Van de Velde, 1993; Clancey 1992). Model instances, sometimes referred to as *case-models*, are state descriptions that contain information about the current state of the world. The role of domain knowledge is to enlarge the case-model, controlled by a suitable method for the task/subtask in question. Characteristic for the view of problem solving as model construction is that the entire case model is considered at each problem-solving step. This is different from searching for just a particular value (i.e. a solution). This view therefore fits very well with the CBR problem solving cycle, where the initial case (the problem to be solved) is an instantiation of the generic case model that becomes enlarged through influence from retrieved cases and adaptation knowledge until it contains a solution that satisfies the initial task requirements. This is also advocated in the framework of Plaza and Arcos (2000). Retainment, i.e. learning, sets up an additional task which takes the final state of the problem solving case and constructs a knowledge structure to be integrated into the case base. Learning can also be viewed as a type of problem solving, i.e. solving a learning problem. This unified view to problem solving and learning also opens up for tightly integrated problem solving and learning architectures (Van de Velde and Aamodt, 1992).

The power of using the three perspectives (tasks, methods, and models) for knowledge level modeling lies in the interaction between the perspectives, and the constraints they impose on each other. For example, a task may be decomposed in two principle ways: By a method-oriented decomposition, or by a model-oriented decomposition. In the former, the type of task decomposition method chosen for the task determines subtasks of a task. For example, a common problem solving method called ”cover-and-differentiate” will decompose a task into two sets of subtasks: One which will try to find solutions that cover for the observations made, and another that tries to differentiate between possible solutions in order to find the best one. In medicine this is known as differential diagnosis, and the method sometimes referred to as the hypothetico-deductive method of reasoning. The method ”hierarchical design” will decompose a design task into a set of course-grained components, which in turn are decomposed into more detailed components, etc. In a model-oriented decomposition, the subtasks of a task are chosen according to what type of domain-models they relate to and the type of case-models they produce. An example would be to decompose a task into a subtask that handles the input of component information, another that deals with process information, etc. The framework opens up for, for example, interrelating cases that cover several subtasks, and to develop problem solving methods and domain models used in retrieval and adaptation suited to each subtask/subcase. A suitable way to integrate knowledge-level and symbol-level modeling will have to be a necessary part of the framework (Aamodt, 1995; Winnem, 1996).

Several issues on the current research agenda of the knowledge modeling community are regarded of particular relevance to CBR knowledge modeling, and form a basis for continued research in our group. First, a lot of work is currently being put into the definition and reuse of ontologies, for domain knowledge as well as task and method knowledge (Gennari et. al., 1994; Schreiber et. al., 1995; Fensel et. al., 2000, Smith et. al. 2006). This will enable the knowledge content of the case vocabulary container, as well as supporting general domain knowledge, to be more easily identified and related to task and method ontologies in a systematic way. In turn this facilitates the sharing of knowledge for CBR systems development (Bergmann et. al., 1997). Second, the understanding of the role of problem solving methods, and their interrelations with tasks and domain knowledge has improved (Aamodt et al., 1992; Benjamins and

Pierret-Golbreich, 1996; Motta and Zdrahal, 1998). The case knowledge as such can be modelled by taking all the three knowledge types into account, and the knowledge-level analysis will help identifying the relationships between the different parts of the case contents.

For CBR as well as other integrated methods, a challenge is the bridging from the knowledge level modeling view to symbol level design and implementation (Leake and Wilson, 1998). Further, although the framework presented is targeted to the modeling and maintenance of application knowledge, the knowledge-level framework also enables the explicit modeling of introspective tasks and their accompanying reasoning methods and reasoning domain models (Ram and Leake, 1995; Cox 2009).

4. A life cycle model of knowledge modeling and learning.

In this section a generic knowledge modeling cycle is presented, as a high-level process model. It is based on the combination of a basically top-down driven, constructive approach to initial knowledge acquisition and modeling, and the bottom-up modeling view represented by continuous learning through retaining problem solving cases as they are experienced.

The objective of the initial knowledge modeling task is to analyze the domain and task in question, to develop the conceptual, mediating models necessary for communication within the development team, and to design and implement the initial operational and fielded version of the system. Initial knowledge modeling, in this sense, covers all phases up to the realization of a computer system according to specifications.

The knowledge maintenance task takes over where the initial knowledge modeling ends, and its objective is to ensure the refinement and updating of the knowledge model as the system is being regularly used. This include to correct errors and improve the knowledge quality, to improve performance efficiency, and to adjust system behaviour according to changes in the surrounding environment, such as changing the type of users interacting with the system or the type of use made of it. The knowledge maintenance task continues throughout the entire lifetime of the system. In Figure 2 (copied from Aamodt, 1995) the two outer, large boxes (with rounded corners) illustrate these two top-level tasks of the knowledge modeling cycle.

More specific modeling subtasks are indicated by grey background. Sharp-cornered boxes are knowledge models, i.e. input and output from the modeling tasks. Arrows indicate the main flow of knowledge and information, and show the most important input/output dependencies between subtasks and models. As shown by the area where the two large boxes overlap, the conceptual knowledge model and the computer internal model are shared by subtasks of both initial knowledge modeling and knowledge maintenance.

The knowledge maintenance task has two subtasks as indicated in the figure. The sustained learning task directly updates the computer internal model each time a new problem has been solved. The other subtask is a periodic and more substantial revision process, i.e. a more thorough analysis, which in this model is assumed to be made after some amount of new experience has been gathered. As illustrated, this revision task may lead directly to the modification of the symbol level model (computer internal model), but it may also go through an update of the knowledge level model (conceptual knowledge model) first.

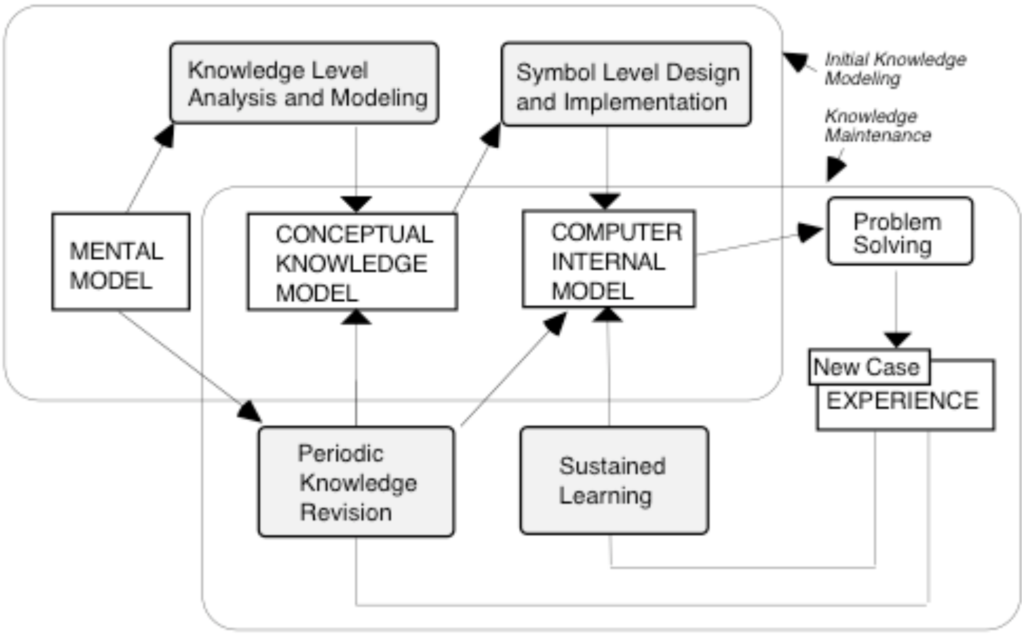

Figure 2 A knowledge modeling and learning cycle

The sustained learning task, on the other hand, regards each new problem solving episode, i.e. each problem just solved, as a source for immediate learning. This implies that the knowledge model (read: the knowledge base) is updated each time a problem is solved.

5. Excerpts from our research

The high level question behind our research within the above framework is: Given the two main types of knowledge, concrete experiences and generalized knowledge, how can they be combined to improve performance beyond that possible by one of them alone?

Historically, this research agenda started with the analysis the lead to the Creek architecture and representation language for knowledge intensive case-based reasoning (Aamodt, 1991). In Creek the knowledge level modeling language and the symbol level representation language has the same high-level structure and visual appearance: A frame-based system, in the original sense of frames as stereotypical (prototypical) concept descriptions. General domain concepts as well as cases are first class concepts, as are also relation types. The core system of interconnected frames may also be viewed as a semantic network. Inference methods are frame matching, value constraint propagation, and a set of inheritance methods. The most advanced of the latter is a method of plausible inheritance in which values can be inherited along several relation types, not only the basic taxonomical relations (Sørmo, 2000). Within Creek, various long-term studies (PhD theses, EU and NFR projects) have been performed. This include the development of a knowledge-level inspired task-oriented architecture for medical image interpretation (Grimnes and Aamodt, 1998), an abductive reasoning approach to context modeling and learning in integrated reasoning systems (Özturk and Aamodt, 1998), and a combined data mining and case-based decision support system in the oil drilling domain (Skalle and Aamodt, 2004). Related to that work a way to combine CBR with probabilistic causal models, i.e. Bayesian Networks, was proposed

(Langseth and Aamodt, 1999). Oil drilling was also the domain for a study of reasoning with time-sequence cases, based on temporal intervals (Jære et. al., 2002).

Some activities explore methods and systems where model-based reasoning plays a strong part in itself, rather than only as part of the CBR process. In one PhD project a method for generating and evaluating explanations for intelligent tutoring systems was proposed (Sørmo and Aamodt, 2002). In another a method for generating explanations for gene–gene relationships and dependencies in order to understand the development of diseases at the level of functional genomics (Kuznierczyk et. al., 2004). This study later shifted focus towards analysis of biomedical ontologies, and particularly the Gene Ontology system (Kuznierczyk, 2004).

More generally, our research into knowledge-intensive case-based explanation studies the combined use of case-specific and general domain knowledge from the perspective of user-targeted explanations (Sørmo et. al., 2005), as well from the perspective of the system-internal explanation methods in CREEK. The transparency of the knowledge representation system in CREEK favours studies of mutual explanation mechanism, i.e. explanation methods serving both purposes (Kofod-Petersen et. al., 2008). This was also studied within a PhD project on conversational case-based reasoning for software component reuse (Gu and Aamodt, 2006-1), where we also focused on how to evaluate CBR systems, through experiments with several evaluation approaches (Gu and Aamodt, 2006-2). Quite another issue was studied in a PhD research done within the EU project Ambiesense (Kofod-Petersen and Aamodt, 2003), where an agent-based architecture was developed for CREEK, aimed to provide contextualized information to mobile users on business or tourist travels.

Additional methods for representing and reasoning with general domain knowledge have been – and are being - being explored. In particular, the studies of Bayesian Networks has given additional insights to the knowledge modeling and representation issues of integrated reasoning system, as well as data mining methods for learning of general domain knowledge. Examples of smaller project that have developed additional demonstrators as part of MSc works, include an ANN system for face recognition integrated into CREEK (Engelsli, 2003), and a text mining system for extracting general domain relationships from text (Thomassen, 2003). Sometimes, it is also useful to lean back and take a look at the more fundamental issues related to developing CBR systems and other AI systems, such as relating current practice to totally different development and modeling views, such as one suggested by an antipoetic analysis (Svedberg, 2005).

More recently we have also studied CBR-specific as well as generalization-based methods triggered by problems that have come out of integrated reasoning studies. This ranges from real-time case-based reasoning for RTS games, implemented in a version of Warcraft [Szczepanski-2009], to the study of ontology foundations and ontological engineering for improving the Gene Ontology, widely used in bioinformatics and genetic engineering. Currently, integrated reasoning underlies on-going work in the domain of petroleum engineering (Shokouhi, 2009), in cooperation with the Department of Geophysics and Petroleum Engineering and the company Verdande Technology. Verdande Technology AS (verdandetechnology.com) was founded a few years ago as a spin-off from our research group at IDI NTNU. Building upon the Creek system, the company builds systems that help oil well drilling engineers to capture cases from interesting situations that occurs during drilling, and retrieve them when needed to interpret new similar situations. In this way the systems can alarm the user if a threatening situation approaches, suggest remedies based on past best practice, and in general provide valuable additional insight into the ongoing drilling process.

We have also ongoing activities in the domains of medicine and in fish farming. In medicine we cooperate with St. Olavs Hospital and NTNU's Department of Molecular Biology and Cancer Research in two project, one funded by EU and one by NFR. Case-based reasoning is a promising technology for a more patient-centered clinical practice than what can be offered from systems that rely on epidemiological studies and statistical variations only. Both projects target patients in palliative care (Kaasa, 2008), with a focus on treatment of pain and depressions, respectively, and an emphasis on improved utilization of clinical guidelines. In aquaculture and fish farming we cooperate with Sintef Fisheries and Aquaculture, under an SFI program granted by NFR. The task is to improved the planning of fish farming operations. In an ongoing PHD study, which is linked to the medical projects and where the focus is on meta-level reasoning and introspective learning, we address the problem of how to utilize a number of different methods in the best way for an input task (Houeland and Aamodt, 2009).

The outcome of these research efforts has been a continually improved understanding of the characteristics and roles for cases in integrated reasoning systems. Our plans for future research include a stronger emphasis on integration of probabilistic models, and a move towards case contents that are unstructured, i.e. natural language text. Both choices are motivated by recent development in these areas internationally as well as high level competence within other AI groups within our Division of Intelligent systems. As for other problem domains it will be interesting to address problems of visual and auditory interpretation and subsequent action, e.g. in robots, by means of integrated CBR systems. Given the versatility of combining situation-specific and generalized knowledge, whether based on symbolic or numeric models, the question is not how to find interesting applications domains. The challenge is rather to limit them to a few that will serve as effective test domains, and also have an impact as eyebrow-lifting demonstrator applications.

References

Aha, D.W., Wettsceherck, D. 1997. Case-based learning: Beyond classification of feature vectors. In Proceeding of ECML-97, European Conference on Machine Learning. Prague, 1997. (www.aic.nrl.navy.mil/~aha/ecml97-wkshp/).

Aha, D.W. 2000. Interactive Case-Based Reasoning: Influences, Utility, and Outlook in an Applied World, International Conference on Industrial and Engineering Applications of Artificial Intelligence & Expert Systems. Keynote talk. New Orleans, LA, 21 June 2000. (www.aic.nrl.navy.mil/~aha/).

Aamodt A.: A Knowledge-Intensive Integrated Approach to Problem Solving and Sustained Learning. PhD. Dissertation. University of Trondheim, Department of Electrical Engineering and Computer Science, Trondheim (1991). [Downloadable from authors publications homepage].

Aamodt, A., Benus, B., Duursma, C., Tomlinson, C., Schrooten, R., and Van de Velde, W. 1992. Task Features and their Use in CommonKADS. KADSII Report Free University of Brussels, VUB/014. vi+122 pgs.

Aamodt, A. 1994. Explanation – driven-case-based reasoning. In Topics in case-based reasoning, edited by S. Wess et al., Springer Verlag, 274-288.

Aamodt, A 1995. Knowledge Acquisition and Learning from Experience - The Role of Case-Specific Knowledge, In Gheorge Tecuci and Yves Kodratoff (eds): Machine learning and knowledge acquisition; Integrated approaches, (Chapter 8), Academic Press, 197-245.

Klaus-Dieter Althoff, Agnar Aamodt. 1996. Relating case-based problem solving and learning methods to task and domain characteristics; towards an analytic framework. AI Communications - The European Journal of Artificial Intelligence, 9 (3):109-116.

Armengol, E., Plaza, E. 1993. Case-based reasoning at the knowledge-level: An analysis of CHEF. In Proceedings of EWCBR-93, First European Workshop on Case-Based Reasoning. University of Kaiserslautern, 290-95.

Benjamins, R., Pierret-Golbreich, C. 1996. Assumptions of problem solving methods. In Proceedings of the 9th European Knowledge Acquisition Workshop, EKAW-96, N. Shadbolt and K. O'Hara and G. Schreiber (eds), 1-16.

Bergmann, R., Wilke, W., Althoff, K.-D., Breen, S. & Johnston, R. 1997. Ingredients for Developing a Case-Based Reasoning Methodology. In: R. Bergmann & W. Wilke (eds.), Proc. of the Fifth German Workshop on Case-Based Reasoning, 49-58

Breuker, J.A., Van de Velde, W. 1994. CommonKADS Library for expertise modelling. IOS Press, Amsterdam.

Clancey, W. J. 1992. Model construction operators. Artificial Intelligence, 53:1-115.

Engelsli, S.E.: Intergration of Neural Networks in Knowledge - Intensive CBR, MSc thesis, Norwegian University of Science and Technology (NTNU), Department of Computer and Information Science, 2003.

Cox, Michael. Metareasoning – a manifesto. Report, BBN Technologies. 2007.

Fensel, D. Horrocks, I., Van Harmelen, F., Decker, S. Erdmann, M. Klein, M. 2000. OIL in a nutshell. In R. Dieng (ed.), In Proceedings of the 12th European Workshop on Knowledge Acquisition, Modeling, and Management (EKAW'00). Lecture Notes in Artificial Intelligence, Springer, 1-16.

Fuchs, B., Mille, A. 1999. A knowledge-level task model of adaptation in case-based reasoning. In Case Based Reasonning Research and Development, Third International Conference on Case-Based Reasoning, Seeon Monastery, Germany, July 1999. Lecture Notes in Artificial Intelligence 1650, Springer-Verlag, 118—131.

Gennari, J. H., Tu, S.W., Rothenfluh, T.E., Musen, M.A. 1994. Mapping domains to methods in support of reuse, International Journal of Human-Computer Studies, 41, 1994, pp 399-424.

Morten Grimnes, Agnar Aamodt: A two layercase-based reasoning architecture for medical image understanding. In Smith, I., Faltings, B. (eds). Advances in case-based reasoning, (Proc. EWCBR-96), Springer Verlag, Lecture Notes in Artificial Intelligence 1168, 1996. pp 164-178.

Mingyang Gu, Agnar Aamodt: Evaluating CBR systems using different data sources: a case study. To appear in Proceedings of the 8th European Conference on Case-Based Reasoning, ECCBR 2006, Öludeniz/Fethiye, Turkey, September 4th-7th, 2006. Lecture Notes in Computer Science, Springer Verlag.

Mingyang Gu, Agnar Aamodt: Dialog Learning in Conversational CBR. Proceedings of the 19th International FLAIRS Conference (Florida Artificial Intelligence Research Society), Melbourne Beach, Florida, May 11-13, 2006, AAAI Press. pp 358-363.

Tor Gunnar Houeland, Agnar Aamodt: Towards an Introspective Architecture for Meta-level Reasoning in Clinical Decision Support Systems. Proceedings of the Workshop on CBR in the Health Sciences, 8th International Conference on Case-Based Reasoning, ICCBR 2009, Seattle, July 2009.

Martha Dørum Jære, Agnar Aamodt, Pål Skalle: Representing temporal knowledge for case-based prediction. Advances in case-basedreasoning; 6th European Conference, ECCBR 2002,Aberdeen, September 2002. Lecture Notes in Artificial Intelligence, LNAI 2416,Springer, pp. 174-188.

Kaasa S, Loge JH, Fayers P, Caraceni A, Strasser F, Hjermstad MJ, Higginson I, Radbruch L, Haugen DF, on behalf of the EPCRC. Symptom assessment in palliative care: A need for international collaboration. J Clin Oncol 2008

Klinker G., Bohla C., Dallemagne G., Marques D. and McDermott J. 1991. Usable and reusable programming constructs. Knowledge Acquisition 2:117-136.

Anders Kofod-Petersen, Agnar Aamodt: Case-based situation assessment in a mobile context-aware system. Proceedings of AIMS2003, Workshop on Artificial Intgelligence for Mobil Systems, Seattle, October, 2003.

Anders Kofod-Petersen, Jörg Cassens, Agnar Aamodt: Explanatory Capabilities in the CREEK Knowledge-Intensive Case-Based Reasoner, Explanatory Capabilities in the CREEK Knowledge-Intensive Case-Based Reasoner. Proceedings of the 10th Scandinavian Conference on Artificial Intelligence (SCAI 2008). Frontiers in Artificial Intelligence and Applications, Volum 173. s. 28-35. IOS Press, 2008.

Waclaw Kusnierczyk, Agnar Aamodt and Astrid Lægreid: Towards Automated Explanation of Gene-Gene Relationships. RECOMB 2004, The Eighth International Conference on Computational Molecular Biology, Poster Presentations, E9, San Diego, March 2004.

Waclaw Kusnierczyk: Taxonomy-based partitioning of the Gene Ontology. *Journal of Biomedical Informatics* 41(2): 282-292 (2008)

Helge Langseth, Agnar Aamodt, Ole Martin Winnem: Learning retrieval knowledge from data. In Sixteenth International Joint Conference on Artificial Intelligence, Workshop ML-5: Automating the Construction of Case-Based Reasoners. Stockholm 1999. Sarabjot Singh Anand, Agnar Aamodt, David W. Aha (eds.). pp. 77-82.

Linster M. 1992. Linking models to make sense to modeling to implement systems in an operational modeling environment. In T. Wetter, K.-D. Althoff, J. Boose, B.R. Gaines, M. Linster, F. Schmalhofer, eds.: Current Developments in Knowledge Acquisition: Proc. of the 6th European Acquisition Workshop EKAW'92, Springer Verlag, 55-74.

Leake D. 1993. Focusing construction and selection of abductive hypotheses. In Proceedings of IJCAI 1993, Chambery, France, Morgan Kaufmann, 24-31.

Leake, D. B., Wilson, D. C. 1998. Categorizing Case-Base Maintenance: Dimensions and Directions. Advances in Case-Based Reasoning: Proceedings of EWCBR-98, Springer-Verlag, Berlin. 13 pgs.

Leake, D. B., Wilson, D. C. 1999. Combining CBR with Interactive Knowledge Acquisition, Manipulation and Reuse. Proceedings of the Third International Conference on Case-Based Reasoning, ICCBR-99, Springer-Verlag, Berlin. 15 pgs.

Motta, E., Zdrahal, Z. 1998. A library of problem-solving components based on the integration of the search paradigm with task and method ontologies. International Journal of Human-Computer Studies 49(4):37-470.

Motta, E., Fensel, D., Gaspari, M., Benjamins, R. 1999. Specifications of Knowledge Components for Reuse. Eleventh International Conference on Software Engineering and Knowledge Engineering (SEKE '99).

Musen, M.A. 1989. Conceptual models of interactive knowledge acquisition tools. Artificial Intelligence 1(1):73-88.

Newell A. 1982. The knowledge level, Artificial Intelligence, 18:87-127.

Plaza, E., Arcos, J.-L. 2000. Towards a Software Architecture for Case-based Reasoning Systems. Foundations of Intelligent Systems, 12th International Symposium, ISMIS 2000. Ras, Z. W. and Ohsuga, S., (Eds.), Lecture Notes in Computer Science, V1932.

Ram, A., Leake, D. 1995. Learning, goals, and learning goals. In A. Ram & D. Leake, Goal-driven learning, MIT press, 1995. Ch.1:1-37.

Richter, M.M. 1995. The knowledge contained in similarity measures. Invited talk at ICCBR-95, Sesimbra, Portugal. (http://wwwagr.informatik.uni-kl.de/~lsa/CBR/Richtericcbr95remarks.html)

Schreiber, A. Th., Wielinga, B. J., Jansweijer, W. H. J. 1995. The KACTUS view on the 'O' word. In IJCAI Workshop on Basic Ontological Issues in Knowledge Sharing, 1995. Also in: J. C. Bioch and Y.-H. Tan (eds.). Proceedings 7th Dutch National Conference on Artificial Intelligence NAIC'95, EURIDIS, Erasmus University Rotterdam, The Netherlands, 159–168.

Samad Valipour Shokouhi, Agnar Aamodt, Pål Skalle and Frode Sørmo: Determining root causes of drilling problems by combining cases and general knowledge . Proceedings of the 8th International Conference on Case-Based Reasoning, ICCBR 2009, Seattle, July 2009. Steels L. 1990. Components of expertise. AI Magazine, 11(2), Summer, 29-49.

Barry Smith, Waclaw Kusnierczyk, Daniel Schober, Werner Ceusters: Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain. KR-MED 2006.

Steels L. 1993. The componential framework and its role in reusability, In J-M. David, J-P. Krivine, R. Simmons (eds.), Second generation expert systems, Springer Verlag, 273-298.

Svedberg, P.: Steps towards an empirically responsible AI; A theoretical and

methodological framework. MSc thesis, Norwegian University of Science and Technology

(NTNU), Department of Computer and Information Science, 2004.

Tomasz Szczepanski, Agnar Aamodt: Case-based Reasoning for Improved Micromanagement in Real-Time Strategy Games. Proceedings of the Workshop on Case-Based Reasoning for Computer Games, 8th International Conference on Case-Based Reasoning, ICCBR 2009, Seattle, July 2009.

Frode Sørmo: Plausible Inheritance; Semantic Network Inference for Case-Based Reasoning. MSc thesis, Norwegian University of Science and Technology (NTNU), Department of Computer and Information Science, 2000.

Frode Sørmo, Agnar Aamodt: Knowledge communication and CBR. 6th European Conference on Case-Based Reasoning, ECCBR 2002, Aberdeen, September 2002. Workshop proceedings. Robert Gordon University, pp. 47-59.

Frode Sørmo, Jörg Cassens, Agnar Aamodt: Explanation in Case-Based Reasoning; Perspectives and Goals . Artificial Intelligence Review. Vol 24, no. 2, October 2005. pp. 109143.

Tomassen, S.L.: Semi-automatic generation of ontologies for knowledge-intensive CBR. MSc thesis, Norwegian University of Science and Technology (NTNU), Department of Computer and Information Science, 2003.

Van de Velde W., Aamodt, A. 1992. Machine Learning Issues in CommonKADS. KADSII Report Free University of Brussels, D 2.11, VUB/002/3.0.

Van de Velde W. 1993. Issues in knowledge level modeling, In J-M. David, J-P. Krivine, R. Simmons (eds.), Second generation expert systems, Springer. pp 211-231.

Wielinga, B., Van de Velde, W., Schreiber, G., Akkermans, H. Towards a unification of knowledge modelling approaches. Proceedings of JKAW-92, Japanese Knowledge Acquisition Workshop. 1992 (17).

Winnem, O.M. 1996, Integrating knowledge-level and symbol-level modelling: The Creest workbench. Master Thesis, Norwegian University of Science and Technology, Department of Informatics. Trondheim, 114 pgs.

Pinar Ozturk, Agnar Aamodt: A context model for knowledge-intensive case-based reasoning. International Journal of Human Computer Studies. Vol. 48, 1998. Academic Press. pp 331-355.