

Towards an Introspective Architecture for Meta-level Reasoning in Clinical Decision Support Systems

Tor Gunnar Houeland and Agnar Aamodt

Department of Computer and Information Science,
Norwegian University of Science and Technology,
NO-7491 Trondheim, Norway
{houeland,agnar}@idi.ntnu.no

Abstract. The paper presents core elements of a meta-level architecture for clinical decision support, in the domain of palliative care. The goal of the reported research is to develop an architecture and an integrated set of methods for an introspective meta-level reasoner. Within the architecture a system is under development that addresses the identification and utilization of clinical guidelines for the assessment and treatment of cancer pain. Case-based reasoning is a core component of the architecture, which also incorporates rule-based and probabilistic model-based methods. The paper presents the overall architectural constraints and exemplifies parts of it through structured component descriptions.

1 Introduction

A clinical decision support system that covers several clinical tasks, such as patient examination, disease hypotetization, diagnosis determination, treatment planning, and drug administration, would typically need to combine several types of knowledge and several reasoning methods to provide good advisory support. There has recently been a renewed interest in meta-level reasoning as a means for a system to effectively combine several reasoning methods [1].

The focus of the research presented here is on meta-level reasoning for clinical decision support. A component-based architecture and an integrated set of methods for a meta-level reasoner to improve its reasoning and learning abilities are currently being developed. The architecture allows for the integration of all three major reasoning paradigms in symbolic AI, i.e. rule-based, case-based, and (deeper) model-based reasoning.

We are addressing this problem in the domain of clinical decision support for palliative care. In a cooperation with the Palliative Medicine Unit, Cancer Department, St. Olavs University Hospital in Trondheim, we are studying the potential for proactive, advice-giving systems to improve treatment of pain in cancer patients. In our current project, short-named TLCPC, the focus is on lung cancer. This research is tightly linked to a larger EU project called EPCRC [2], which covers all forms of cancer pain as well as other problems related to

palliative care for long-term cancer patients. A motivation for that project is also to standardize procedures and to unify clinical practice in palliative care [3], a goal for which computerized decision support systems have a strong potential. The system under development will in particular address the identification and utilization of clinical guidelines.

In the next section we review some of the related research in metareasoning with case-based components, and CBR used for guideline supported clinical decision-making. In section 3 our metareasoning approach is introduced. Essential components of the architecture are discussed in section 4. The status of the clinical guideline application is presented in section 5. Concluding remarks end the paper.

2 Related Research

There is a significant amount of research that has addressed metareasoning in relation to CBR systems. In Meta-AQUA [4, 5] introspection is used in the retain phase to learn from mistakes. What is referred to as *introspective meta-XPs* are used to represent failures encountered while the system operates. The system constructs learning plans that consist of calling various learning algorithms.

Early accounts of meta-level architectures involving CBR also include BOLERO [6] and ANALOG [7]. In BOLERO, a case-based meta-level planner controls the execution of a rule-based reasoner designed to accomplish different medical diagnosis tasks. New plans are constructed during the problem solving process when needed, and captured as new cases by the meta-level learner. In ANALOG a selected method runs until an impasse situation is encountered, at which time a new method selection process is run. The meta-level learner remembers failed and successful instantiations of methods.

Christodoulou and Keravnou describe a meta-level architecture [8] in the domain of breast cancer histopathology. Each problem solver is associated with a task, an inference mechanism, and domain knowledge constraints. A set of meta-parameters is used by the metareasoner to characterize knowledge types (e.g. experience-based, causal), and desired solution properties (e.g. level of detail, accuracy, efficiency). The metareasoner is a case-based reasoner that captures and stores problem-solving paths as strategy cases incrementally.

In the ADAPtER system [9], an architecture that combines model-based and case-based reasoning for medical diagnosis, the CBR component is the primary object level reasoner. The model-based method is triggered either when the retrieval method fails to find a matching case, or the case adaptation module fails. Cases are captured as a compilation of the model-based process.

The principles of evidence-based medicine, in which systematic research evaluation regimes are used to assess and justify results from medical research, form a well-established basis for medical practice [10]. One manifestation of this is the specification of clinical guidelines, i.e. operational procedures for how to conduct a particular type of patient examination, make a diagnosis, administer a type of drug, or perform other types of treatment. Another source of information used

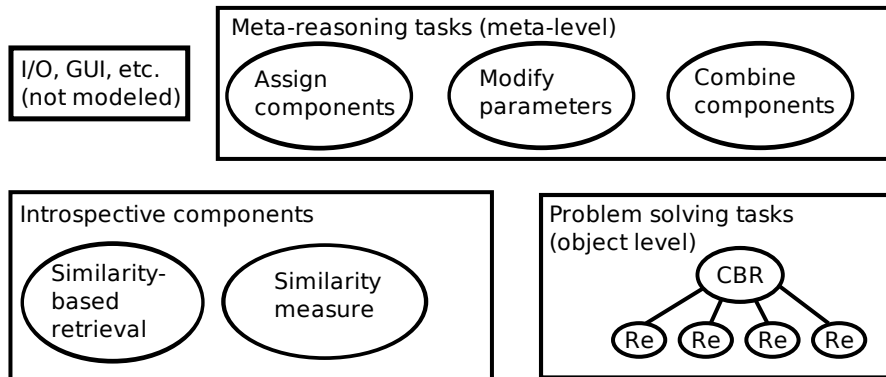


Fig. 1. A CBR system based on the introspective architecture

by clinicians in their daily practice is, of course, the set of experiences a clinician has from earlier patients. Hence, a combination of general guidelines with past experience cases is potentially a strong combination. Identifying the strengths and weaknesses of a case-based method vs. a generalization-based one can be done analytically or experimentally. Marling et. al. [11] reports on an experiment in the domain of nutritional menu planning, in which a hybrid system was developed by combining the strengths of separate rule-based and case-based reasoners. The hybrid system outperformed both separate systems.

In the CARE-PARTNER system [12] medical guidelines are realized in the form of problem solving pathways, implemented as a rule-based system combined with information retrieval, and with CBR as the main problem solving method. GLARE [13] is a general decision-support system for managing and utilizing clinical guidelines, in which guidelines are represented as a hierarchical structure of decision paths. A CBR system was incorporated as additional knowledge to handle situations that are not covered by the guidelines, so-called non-compliances [14]. Even if several medical guideline systems combine different reasoning paradigms, there has been very little work on moving the combined reasoning up to the metareasoning level.

3 Metareasoning Approach

Recently, Cox and Raja [1] presented a general high-level framework of metareasoning. It relates three different levels, the ground level (physical perception and action), the object level (reasoning about action), and the meta-level (reasoning about reasoning). Leake and Wilson [15] address the learning part of metareasoning within this framework, and present a set of challenging issues, such as learning for self-understanding and self-explanation.

A system using our introspective architecture with CBR reasoning methods which follows this general framework is shown in figure 1. The ground level is represented as I/O for the system, which is not covered in our architecture but

must necessarily exist in some form for a complete CBR system. In the shown system reasoning tasks are explicitly represented, with a “CBR-based problem solving task” at the object level and several metareasoning tasks.

In our architecture there is no fundamental difference in how reasoning tasks for the two levels must be represented, but they are shown as separate boxes in the figure to ease understanding and fit the metareasoning framework of Cox and Raja. The importance of our introspective architecture lies in the components, each of which implements a specific functionality, and is modified and assigned to a particular reasoning task by meta-level reasoning components.

Many meta-level reasoning systems start by building a partial solution and then invoking meta-level reasoning functionality to determine whether the object-level reasoning processes are working satisfactorily or not. Such systems then either proceed as normal, or classify the reasoning process as a failure and create a new metareasoning goal to learn from this failure.

This is in contrast to the meta-level control agent in our approach, which operates on the object-level reasoning components directly without a specific reasoning failure to address. This allows the system to have a clearer broad focus on performing changes that affect the entire reasoning system instead of correcting single failures, and this broadening is also identified by Leake and Wilson [15] as an important opportunity for a more flexible learning focus. Future planned meta-level components in our approach such as a competence-evaluator for problem-solving methods can also assist in providing the system with a level of self-understanding, which is another identified opportunity for introspective learning systems.

However, this illustrates an interesting trade-off compared to other approaches. In our approach the meta-level components are explicitly meant to affect the system’s reasoning processes, and the architectural design has been created to make this feasible and as easy as possible. But we do not have a so-called narrow focus on repairing specific portions of the system [15], and in practice this means that our approach is not currently as well developed for making these repairs as some other systems. Due to our focus we do not yet have a meta-level component dedicated to detecting and repairing failures during object-level reasoning, and as such the individual problem solving processes are only improved indirectly by selecting the best available components. Unlike several other systems, if this competence evaluation is incorrect it will not be detected by our current set of meta-level components during the problem solving process itself, and the lessons learned from this failure will only be available for subsequent problem solving instances.

Richter [16] introduced the knowledge container model, where pieces of knowledge can be categorized into four categories: case vocabulary, cases, similarity assessment and adaptation knowledge. Given that a set of precise and consistent terms is important for reasoning at the meta-level, we explicitly include the CBR system vocabulary as part of the case vocabulary category.

To make sure that the vocabulary contains semantics that are useful for human designers, we suggest that the vocabulary should be created and updated

manually, unlike the other knowledge sources which will be changed and influenced by the meta-level reasoner. In our architecture the case knowledge is stored in cases, similarity assessment as part of the components providing methods for object-level retrieval and the adaptation knowledge is shared between object level reuse components and a more profound system-level adaptation in meta-level components.

For the meta-level control agent to be able to examine and calibrate the system's components, it's important that the settings available can be interpreted by the control program, and that the components expose interfaces at an appropriate level of abstraction. To facilitate this, we are developing a vocabulary of CBR-related terms and the intended semantics as part of our approach.

An important aspect of the vocabulary is that it describes what the terms mean at a semantic level without relying on the specific realizations in any particular CBR system implementation. An example of this is a case base, which is simply defined as a set of cases. While in practice many CBR systems store the cases as a form of ordered lists, their reasoning processes typically do not rely on the particular order the cases are listed in within computer memory or on disk. The purpose of the vocabulary is to abstract away the particular specifics in implementations and generalize the terms to a semantic level where it describes the actual requirements for a term without imposing a needlessly specific design.

Our vocabulary is being developed to be conceptually compatible with the terms used in CBR_{Onto} [17], a modeling framework that has already examined the meanings and relations between terms from an ontological perspective. As an added benefit the CBR_{Onto} ontology is compatible with the "4 REs" CBR process cycle [18] which is widely quoted and referred to in the CBR literature.

4 Introspective Architecture

A particularly important part of our semantic task description is the type system for component inputs and outputs. These types represent at the knowledge level the essential meaningful content that is to be processed by each component.

In our conceptual framework we consider the traditional core case-based reasoning process as one possible problem solving method at the object level. We consider this to be one part of the combined reasoning system, with specific responsibilities, and there can be potentially many other problem solving methods implemented in the same architecture.

A high-level task decomposition of the CBR process for our architecture is shown in figure 2, where the task "Case-based reasoning" is split into the 4 RE subtasks, and each of these are further split into smaller subtasks. Any components that are assigned to parts of this process must match the semantic input and output types for the corresponding tasks, which are also included in the figure.

To elaborate on the case-based retrieval task, it starts from a **Problem-Input** and retrieves a set of cases **Set** $\langle CB \rangle$ where CB is a type adhering to the **Case** semantics, and is further subdivided into subtasks that further specify

Task	Input	Output
Case-based reasoning	ProblemInput	ProblemSolution
- Retrieve	ProblemInput	Set<CB>
- Problem characterization	ProblemInput	<i>PC</i>
- Case retrieval	<i>PC</i>	Set<CB>
- Focus	Set<CB>	Set<CB>
- Reuse	<i>(PC, Set<CB>)</i>	ProblemSolution
- Adapt solution method	<i>(PC, Set<CB>)</i>	SolutionMethod
- Adapt solution	<i>(PC, SolutionMethod, Set<CB>)</i>	ProblemSolution
- Revise	ProblemSolution	(PS, SE)
- Evaluate solution	ProblemSolution	SolutionEvaluation
- Repair solution	(PS, SE)	ProblemSolution
- Retain	(PS, SE)	
- Update general knowledge	(PS, SE)	
- Add to case base	(PS, SE)	

(PS, SE) is an abbreviation for **(ProblemSolution, SolutionEvaluation)** due to space concerns.

Fig. 2. Inputs and output for the CBR tasks

how this is performed. It starts with a subtask identifying the important aspects of the problem and characterizing it as a query, transforming the **Problem-Input** into an intermediate form *PC*. This intermediate form is not restricted by the architecture, and the only requirement is that the methods performing the subsequent CBR subtasks can accept the characterization form *PC* as input.

This characterization is used to retrieve a number of previous cases from the case base, and then this is further narrowed to just focus on the most relevant information by e.g. filtering out a subset of cases or generalizing cases. The other subtasks are formalized in a similar way.

In our approach, the metareasoning components exist separately from the object-level CBR reasoning components, and in fact influence and controls how the CBR problem solving method is performed, which corresponds to the aforementioned metareasoning cycle [1].

By evaluating how the CBR method performs while solving actual new problem instances, the meta-level control agent can identify the strengths and weaknesses of the current system and attempt to use this to improve the system's competence or use alternative reasoning methods. For our meta-level component-combiner this is achieved by attempting to re-solve problems using different assignments of methods for each of the tasks and subtasks in the architecture. Whether the the newly combined reasoning process is an improvement is then evaluated based on whether the solutions produced for individual problem queries are correct for more problem instances than before.

4.1 Architectural components

One of the most important features of this architecture is that each component contains extra structures that semantically describe the component semantically using our CBR vocabulary. It is on the basis of this added information that the meta-level control component can automatically assign components to perform the system's reasoning tasks. Figure 3 shows such a self-describing semantic structure for an example object-level retrieval component.

Each component contains a list of types that apply for the component, which is listed first in the figures. This can either just be a single *name* to indicate that any type is supported, or a statement of the form "*name isa supertype*", which indicates that the *name* type must support the same operations as *supertype*. This is useful when a component performs a generic operation that can apply to many different types of input, and e.g. only requires that the input and output types are the same or that two inputs are comparable. This section is empty for simpler components that only refer directly to specific types in the vocabulary, such as the illustrated example retrieval component.

After that the input and output variables are listed. Each variable consists of a line of the format "*name: type*", where *name* is an identifier that is used to refer to the input or output throughout the component specification and *type* is the semantic type, which can either be a specific from the vocabulary or one of the types specified in the component's Types section. An example of this is the line "*query: AttributeCase*" from the example component, which means that the component receives an input of type **AttributeCase** which is referred to as *query* in the component description.

The following section lists Conditions that have to be fulfilled for the component to produce the expected results. As long as the input variables conform to the specified conditions, the output variables are guaranteed to follow the specifications in the Guarantees section. While the type specifications are semantic restrictions on what kind of operation the component can perform, the conditions specify for which values of inputs the component will actually behave as intended, and these conditions are not necessarily checked by the component.

Because of this the component's operation can usually be performed for non-conforming inputs, but this can produce undefined results and should be avoided. By listing the conditions in the self-describing structure, the metareasoning component can make sure that only compatible components are combined, or that the conditions are checked on-demand before the operation is performed where this cannot be guaranteed.

The final section is a short semantic description of the core functionality performed by the component. This has the format "*x based on y*", where *x* and *y* are statements composed using the input and output variables as well as a set of pre-specified terms representing important concepts related to the reasoning system and the application domain.

By combining the example component in figure 3 with a component providing similarity assessments, the meta-level control agent can create a new component that accepts a query and casebase of type **AttributeCase** and returns a ranked

Types:
Input: <i>query</i> : AttributeCase <i>casebase</i> : Set<AttributeCase> <i>similarityfunction</i> : (AttributeCase <i>source</i> , AttributeCase <i>target</i>) → Similarity
Output: <i>ordered</i> : List<AttributeCase>
Conditions: $5 \leq \text{sizeof}(\text{casebase}) < 100000$ $0 \leq \text{similarityfunction}(x, y) \leq 1$
Guarantees: $\text{sizeof}(\text{ordered}) = 5$
Approach: $\text{order}(\text{casebase})$ based on <i>query</i>

Fig. 3. An example component that orders a case base based on similarity to a query using a specified similarity function.

list of matching cases. This is done by verifying that it fulfills the semantic type requirements, and any conditions that cannot be guaranteed to be always be fulfilled can be verified at run-time in an implemented system.

If the similarity measure is based on e.g. feature weights, the meta-level learning component can also further refine this into a component that learns the feature weights automatically while it is being used. This can be done by matching it with an appropriate learning method that takes a case base (**Set<AttributeCase>**) as input and produces a similarity importance value for each feature among cases in the casebase.

5 Palliative Care Application

Although there has been a lot of focus on developing clinical guidelines within the medical community, and several guideline systems have been developed by medical professional organizations, there is not a consensus as to what is a good guideline system. Further, the active use of guidelines in a clinical setting is far from the level desired, both from the perspective of quality of treatment and the perspective of unified treatment across hospitals and countries.

In our partner project EPCRC, being a collaboration involving many highly influential medical groups in across Europe, the aim is to reach a consensus on a set of high-quality and operational guidelines that will be used in practice. While awaiting the results from EPCRC, we currently work with an existing set of guidelines defined by NCCN (National Comprehensive Cancer Network) in the US. An ontology is currently being built based on a combination of generic UMLS terms combined with terms from the SNOMED and NCI (National Cancer Institute in the US) ontologies.

The top-down design process is combined with bottom-up experimental system building, starting from simple system components that will be combined. Currently a simple a rule-based reasoner is being implemented at the object level. Example guidelines link patient data related to pain level, pain history, and history of treatment, to the next treatment. Type of treatment considered

is the administration of different types of analgesics, with opioids as the largest subclass. Based on the results from the initial model, the system will be extended with case-based and model-based components. The model-based component will be a Bayesian network for reasoning about causality under uncertainty.

Acquiring the necessary medical knowledge needed for our experiments is a continuing process. To advance the method development we are developing a toy example system in the domain of advice-giving for film selection, in parallel to the more complex clinical guidelines system. In that system CBR methods are currently focused on both the meta and object levels. Cases, problems and solutions are currently represented as feature vectors, and there are no explicitly represented general knowledge structures outside of the CBR reasoning components. To predict a movie rating for a user, a retrieval component retrieves a number of similar other users that have seen the movie, where similarity is determined as the average difference in ratings for movies both users have rated. A simple reuse component then copies the majority rating among the retrieved cases. The system has a meta-level control agent that adheres to the principles of the metareasoning approach described earlier. Component combinations are tried out based on matching the input and output type descriptions. Although simple and different from the clinical application, the system assists in the bottom-up specification of the introspective architecture by providing a test-bed for experiments.

In a clinical guideline support system past cases may be utilized in several ways. The role we have intended for the cases in our guideline system is twofold. Having arrived at a leaf node in the guideline structure, CBR will be used to continue from there by providing a more specific and detailed advice, based on adapting a past result. On the other hand, if the guideline system cannot provide reasonable advice, CBR is triggered as a complementary method. The latter approach is similar to the non-compliance method [14] referred to earlier.

6 Conclusions

In this paper we have presented an introspective approach to meta-level learning and outlined a component-based architecture for designing reasoning systems which supports our introspective methods. The main contribution of our approach is the way our architecture allows for gradual additions of metareasoning methods that focus on broad, system-wide improvements.

We are working on further developing this architecture, and adding new components directed towards both object-level and meta-level reasoning methods for a clinical decision support system based on medical treatment guidelines.

Acknowledgements

This research is supported by Norwegian Research Council (NFR) grant, Contract no 183362, Translational Research in Lung Cancer and Palliative Care (TLCPC), together with funds from the Norwegian University of Science and

Technology (NTNU). Thanks to Tore Bruland and Helge Langseth who contributed on the AI method side through a series of discussions, Sunil Raja together with Heidi Knobel who have provided insight into clinical pain assessment and treatment, and Stein Kaasa who is the initiator and project leader of both the TLCPC and EPCRC projects.

References

1. Cox, M.T., Raja, A.: Metareasoning: A manifesto. Technical report, BBN TM-2028, BBN Technologies (2007)
2. Haugen, D.F., Kaasa, S.: The EPCRC. *Eur J Palliat Care* (2007; 14(3):130)
3. Kaasa, S.: Palliative care research – time to intensify international collaboration. *Palliat Med* (2008; 22:301-2.)
4. Cox, M.T., Eiselt, K., Kolodner, J., Nersessian, N., Recker, M., Simon, T.: Introspective multistrategy learning: On the construction of learning strategies. *Artificial Intelligence* **112** (1999) 1–55
5. Cox, M.T.: Multistrategy learning with introspective meta-explanations. In: *Machine Learning: Proceedings of the Ninth International Conference*, Morgan Kaufmann (1992) 123–128
6. Lopez, B., Plaza, E.: Case-based planning for medical diagnosis. In Ras., Z. (ed.), *Proceedings of ISMIS-93*. LNAI 837, Springer, 1993: 96-105
7. Arcos, J., Plaza, E.: A reflective architecture for integrated memory-based learning and reasoning. In Weiss, S. et. al. (eds.): LNAI 873 , Springer, 1994: 289-300
8. Christodoulou, E., Keravnou, E.: Metareasoning and meta-level learning in a hybrid knowledge-based architecture. *Artificial Intelligence in Medicine* 14 (1998): 53-81
9. Torasso, P.: Multiple representations and multi-modal reasoning in medical diagnostic systems. *Artificial Intelligence in Medicine*, 23 (2001): 49-69.
10. Friedland, D.: Evidence-based medicine: A framework for clinical practice. (1998)
11. Marling, C., Petot, G., Sterling, L.: Integrating case-based and rule-based reasoning to meet multiple design constraints. *Comp. Intell.*, 15 (3), 1999, 308-332
12. Bichindaritz, I., Kansu, E., Sullivan, K.M.: Case-Based Reasoning in CARE-PARTNER: Gathering Evidence for Evidence-Based Medical Practice. EWCBR'98, LNAI 1488, pp. 334-345 (1998)
13. Terenziani, P., Molino, G., Torchio, M.: A modular approach for representing and executing clinical guidelines. *Artificial Intelligence in Medicine*, 23(3):249–276 (2001)
14. S. Montani, S.: Case-based reasoning for managing non-compliance with clinical guidelines. ICCBR 2007, 5th Workshop on CBR in the Health Sciences (2007)
15. Leake, D., Wilson, M.: Extending introspective learning from self-models. *Proceedings of the AAAI 2008 Workshop on Metareasoning: Thinking About Thinking* (2008)
16. Richter, M.M.: The knowledge contained in similarity measures. Invited Talk at ICCBR-95 (1995)
17. Díaz-Agudo, B., González-Calero, P.A.: CBRonto: A Task/Method Ontology for CBR. In: *Procs. of the 15th International FLAIRS'02 Conference*, AAAI Press (2002) 101–105
18. Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications* **7**(1) (March 1994) 39–59