Developing Digital Support for Learning and Diagnostic Reasoning in Clinical Practice

Chunli Yan

Ph.D. Thesis

DEPARTMENT OF COMPUTING SCIENCE
UMEÅ UNIVERSITY
SWEDEN
2018

Department of Computing Science
Umeå University
SE-901 87 Umeå, Sweden

chunli@cs.umu.se

ISBN 978-91-7601-918-4
ISSN 0348-0542
UMINF 18.09

Printed by Print & Media, Umeå University, 2018

Abstract

The two main purposes of clinical decision-support systems (CDSSs) are to provide healthcare professionals decision-making support based on evidence-based medical knowledge, and a continuing medical education. This thesis focuses on both purposes and shows how fundamental theory in the field of artificial intelligence can be developed, adapted and implemented in a CDSS for supporting learning and diagnostic reasoning in clinical practice. The main research problems addressed in this thesis are how to represent and manage uncertain, incomplete, inconsistent and distributed knowledge in automated reasoning and decision-making with the clinicians in the loop, how to facilitate the knowledge engineering and maintenance process, and how to detect and support learning and skill development in CDSS users.

Research contributions include theories, methods, and algorithms based on possibilistic logic and formal argumentation for representing and managing uncertain, incomplete, inconsistent and distributed medical knowledge, and for supporting reasoning and decision-making when using a CDSS. The clinician is provided potentially conflicting arguments and their strength based on different diagnostic criteria and the available patient information in order to make an informed decision. The theoretical results were implemented in the Dementia Diagnosis and Management Support System - Web version (DMSS-W), in a multi-agent hypothesis-driven inquiry dialogue system, and in an inference engine serving as a module of ACKTUS.

CDSS maintenance is challenging since new knowledge about diseases and treatments are continuously developed. Typically, knowledge and software engineers are needed to bridge medical experts and CDSSs, leading to time-consuming system development. ACKTUS (Activity-Centered Knowledge and Interaction Tailored to Users) was, as part of this research, further developed as a generic web-based platform for knowledge management and end-user development of CDSSs. It includes the inference engine and a content management system that the medical expert can use to manage knowledge, design and evaluate CDSSs. A graphical user interface generator synchronizes the interface to the ontology serving as the knowledge base. ACKTUS was used for developing DMSS-W, and facilitated the system development and maintenance.

To offer person-tailored support for the clinician's learning, reasoning and decision-making, the CDSS design was based on theories of how novices and experts reason and make decisions. Pilot case studies involving physicians with different levels of expertise who applied DMSS-W in patient cases were conducted in clinical practice to explore methods for detecting skill levels and whether learning is taking place. The results indicated that the skill levels can be detected using the method. The novice was seen to develop reasoning strategies similar to an expert's, indicating that learning was taking place. In future work, tailored educational support will be developed, and evaluated using the methods.

Sammanfattning

Kliniska beslutsstödsystem är datorsystem som stödjer hälso- och sjukvårdspersonal i beslutsfattande och fortbildning. De två huvudsakliga syftena med kliniska beslutsstödsystem är att sprida evidens-baserad medicinsk kunskap till den kliniska vardagen och att ge en kontinuerlig fortbildning till användarna av systemen. Denna avhandling behandlar båda syftena och bidrar till olika aspekter av utveckling och implementation av kliniska beslutssstödsystem, såsom underliggande teorier för formalisering och exekvering av kunskap, design, struktur, utveckling, implementation och utvärdering. De huvudsakliga forskningsproblem som adresseras är hur representera och hantera osäker, inkomplett, tvetydig och distribuerad information i automatiserat resonerande och beslutsfattande med de kliniskt professionella med i loopen; hur underlägga kunskapsmodellering och underhåll; och hur detektera och stödja lärande och färdighetsutveckling hos användarna av ett kliniskt beslutsstödsystem.

Forskningsresultaten innefattar teorier, metoder och algoritmer baserade på possibilistisk logik och argumentation, som implementerats i ett multiagentsystem för att representera och hantera osäker, inkomplett, tvetydig och distribuerad medicinsk kunskap, och för att stödja resonerande och beslutsfattande vid användning av ett kliniskt beslutsstödsystem. Ett hypotesdrivet "inquiry"-dialogsystem baserad på possibilistisk logik utvecklades för att hantera osäker och konfliktande medicinsk information. Klinikern ges en översikt av argument som kan vara i konflikt, baserade på olika diagnostiska kriterier och den tillgängliga patientinformationen, och kan fatta ett underbyggt beslut utifrån informationen. De teoretiska resultaten implementerades i det kliniska beslutsstödsystemet DMSS-W (Dementia Management Support System-Web version).

Underhåll av kliniska beslutsstödsystem är en utmaning eftersom den medicinska kunskapen om sjukdomar och deras behandlingar ständigt utvecklas. Vanligtvis behöver kunskapsmodelleringsingenjörer fungera som broar mellan medicinska experter och systemen vilket är opraktiskt och tidskrävande. ACKTUS (Activity-Centered Knowledge and Interaction Tailored to Users) vidareutvecklades inom ramen för denna forskning som en generell plattform för kunskapsmodellering och slutanvändareutveckling av kliniska beslutsstödsystem. ACKTUS inkluderar ett innehållsmodelleringssystem som den kliniska experten kan använda för kunskapsmodellering, design och utvärdering av kliniska beslutsstödsystem. En gränssnittsgenerator synkroniserar användargränssnittet med ontologin som tjänar som kunskapsbas. ACKTUS kan förbättra systemutveckling och underhåll, och användes för utveckling av DMSS-W.

Design av det teoretiska ramverket, algoritmer och interaktionen med systemet

är baserade på teorier om hur noviser och experter resonerar och fattar diagnostiska beslut. Detta i syfte att erbjuda klinikern personanpassat stöd för lärande, resonerande och beslutsfattande. Pilotstudier genomfördes i kliniska praktik för att undersöka metoder för att detektera färdighetsnivåer och ifall lärande sker. Användarna av DMSS-W var kategoriserade utifrån expertis genom analys av användardata. Resultatet från fallstudierna indikerade att metoden kan användas för att detektera färdighetsnivåer. Novisen observerades också utveckla resonerandestrategier liknande en experts, vilket indikerade att lärande ägde rum. I framtida studier kommer ett personanpassat stöd för lärande att utvecklas, vilket kommer att utvärderas med hjälp av metoderna.

Sammanfattningsvis visar avhandlingen hur grundläggande teoretiska resultat inom AI-området kan utvecklas, anpassas och implementeras i ett beslutsstödsystem för stöd till lärande och diagnostiskt resonerande i klinisk praktik.

Preface

This Ph.D. thesis consists of the following papers:

Paper I **C. Yan** and H. Lindgren. “Hypothesis-Driven Agent Dialogues for Dementia Assessment”. In Proceedings of VIII Workshop on Agents Applied in Health Care (A2HC) 2013, Murcia, Spain, pp. 13-24, 2013.

Paper II **C. Yan**, H. Lindgren and J.C. Nieves. “A Dialogue-Based Approach for Dealing with Uncertain and Conflicting Information in Medical Diagnosis”. Accepted by the journal Autonomous Agents and Multi-Agent Systems, 2018.

Paper III H. Lindgren and **C. Yan**. “ACKTUS - A Platform for Developing Personalized Support Systems in the Health Domain”. In Proceedings of the 5th International Conference on Digital Health 2015, Florence, Italy, pp. 135-142, 2015.

Paper IV **C. Yan** and H. Lindgren. “A Generic Approach for End-User Development of Clinical Decision Support Systems”. Technical report / UMINF 18.08, ISSN 0348-0542, Umeå University, Umeå, 2018.

Paper V H. Lindgren and **C. Yan**. “Detecting Learning and Reasoning Patterns in a CDSS for Dementia Investigation”. Studies in Health Technology and Informatics, vol 210, pp. 739-742, 2015[1].

Paper VI **C. Yan** and H. Lindgren. “Diagnostic Reasoning Guided by a Decision-Support System: a Case Study”. In Proceedings of the ACM European Conference on Cognitive Ergonomics (ECCE-17), Umeå, Sweden, pp. 25-30, 2017[2].

In addition to the papers included in this thesis, other publications were published within the studies but not contained in this Ph.D. thesis, as follows:

- H. Lindgren, M.H. Lu, Y. Hong and **C. Yan**. “Applying the Zone of Proximal Development When Evaluating Clinical Decision Support Systems: a Case S-tudy”. Studies in Health Technology and Informatics, Vol 247, pp. 131-135, 2018.

[1] Reprinted with permission from IOS press. The publication is available at IOS Press through http://dx.doi.org/10.3233/978-1-61499-512-8-739

[2] Reprinted with permission from ACM.

- J. Baskar, **C. Yan** and H. Lindgren. "Instrument-Oriented Approach to Detecting and Representing Human Activity for Supporting Executive Functions and Learning". In Proceedings of the ACM European Conference on Cognitive Ergonomics (ECCE-17), Umeå, Sweden, pp. 105-112, 2017.
- H. Lindgren, J. Baskar, E. Guerrero, J.C. Nieves, I. Nilsson and **C. Yan**. "Computer-Supported Assessment for Tailoring Assistive Technology". In Proceedings of the 6th ACM International Conference on Digital Health 2016, Montreal, Canada, pp. 1-10, 2016.
- **C. Yan**, J.C. Nieves and H. Lindgren. "A Multi-Agent System for Nested Inquiry Dialogues". In: Y. Demazeau, F. Zambonelli, J.M. Corchado, J. Bajo (eds.) Advances in Practical Applications of Heterogeneous Multi-Agent Systems. The PAAMS Collection. PAAMS 2014. Lecture Notes in Computer Science, vol. 8473, pp. 303-314. Springer, Cham, 2014.
- J. Baskar, H. Lindgren and **C. Yan**. "User's Control of Personalised Intelligent Environments Supporting Health". In Proceedings of the 9th International Conference on Intelligent Environments, Athens, Greece, IEEE Computer Society Press, pp. 270-273, 2013.
- H. Lindgren, F. Yekeh, J. Baskar and **C. Yan**. "Agent-Supported Assessment for Personalized Ambient Assisted Living". In Proceedings of Workshop on Agents Applied in Health Care(A2HC), Valencia, Spain, pp. 141-150, 2012.
- J. Baskar, H. Lindgren, D. Surie, **C. Yan** and F Yekeh. "Personalisation and User Models for Support in Daily Living". In the 27th Annual Workshop of the Swedish Artificial Intelligence Society (SAIS), Örebro, Sweden, pp. 7-15, 2012.
- H. Lindgren, P.J. Winnberg and **C. Yan**. "Collaborative Development of Knowledge-Based Support Systems: a Case Study". Studies in Health Technology and Informatics, vol 180, pp. 1111-1113, 2012.

Acknowledgements

Time flies so fast and I have been already working in *User Interaction and Knowledge Modelling Group* in Umeå University for more than eight years! First of all, I am sincerely thankful to **Helena Lindgren**, who accepted me in your group first as a research engineer and later as your PhD student! Without you, this thesis work would be impossible. You have always being patient, inspiring and supportive in guiding me through my PhD research. You are not only a great supervisor offering invaluable guidance, astute criticism and meticulous suggestions, but also a close friend who helps me with many aspects in daily life in Sweden.

I also would like to thank my co-supervisors, **Juan Carlos Nieves** and **Lars-Erik Janlert**, especially Juan Carlos, who led me to the beautiful world of computational logic and taught me a lot of valuable knowledge.

I wish to express my gratitude to my officemate and roommate, **Jayalakshmi**. Our trips both for work and for leisure time are the good memories that I will never forget. Also, to **Esteban**, thanks for the valuable ideas and for picking up berries together.

I would like to thank all the other group members (previous and current), **Madeleine**, **Johannes**, **Linus**, **Ming-Hsi**, **Timotheus**, **Nazanin**, and **Monika**. Working with you guys made me really happy!

Thank you to **Frank**, **Thomas**, **Kai-Florian** and **Lili** for your valuable advice to my thesis and presentation!

I would also like to take this opportunity to extend my gratitude to everyone at the Department of Computing Science for your support from any side, and in any form.

Last, but not least, I would like to express my immense gratitude to my family. My lovely son and daughter, you colored my gray PhD life and gave me courage and energy when I encountered problems during my PhD study. My dearest lover, my best friend, my husband, **Jie Sun**, you not only took care of most housework, took care of the kids to let me focus on my research, but also helped me a lot with my writing skills. Even more importantly, you encouraged me to conquer difficulties. Apart from my supervisors and myself, you are the one to make this thesis work possible!

Göteborg, August 2018
Chunli Yan

Contents

Chapter 1

Introduction

This thesis provides theory, methods, and algorithms for developing digital support for diagnostic reasoning and learning in clinical practice. This is typically accomplished using clinical decision-support systems (CDSSs). According to Sim et al. [103], CDSSs are defined as software that is:

"...designed to be a direct aid to clinical decision-making, in which the characteristics of an individual patient are matched to a computerized clinical knowledge base and patient-specific assessments or recommendations are then presented to the clinician or the patient for a decision."

The main purposes of CDSSs are to disseminate evidence-based medicine (EBM) and best practice knowledge regarding healthcare and to provide continuing medical education to the users of the CDSS [57]. This thesis addresses both purposes. CDSSs provide medical clinicians with patient-specific diagnostic information interpreted using evidence-based knowledge to enhance healthcare [10, 34, 64, 83, 92]. They assist clinicians with information to support their decision-making, reduce their workload and improve clinical workflows [34]. CDSSs have been shown to improve medical clinicians' performance and education [40]. Consequently, CDSSs are expected to significantly improve healthcare quality, increase healthcare efficiency, reduce cost, and educate users.

This thesis covers several aspects of CDSSs, including the design, structure, development, and the underlying theory and knowledge. This is done by combining and further developing theories and methods in different research domains from (1) a *computational perspective* on how uncertain, incomplete, inconsistent, and distributed information can be represented and managed in a system aimed at supporting human reasoning and decision-making (Chapter 2); (2) a *software engineering perspective* on how to translate domain knowledge expressed in natural language into computer-interpretable formats, and on how to conduct the knowledge maintenance (i.e. keep updated with new knowledge, new disease, new organizations) by end-user development of the CDSS, but with minimal involvement of knowledge engineers and software engineers (Chapter 3); and (3) a *human perspective* on how clinicians reason, make clinical decisions and develop skills, as well as how to provide tailored support for learning, reasoning, and decision-making to an individual medical clinician (Chapter 4).

The following three research questions have been identified and addressed in this thesis. They are complementary and interrelated, and each covers a range of more specific problems and research questions that are addressed in the thesis.

- **How can uncertain, incomplete, inconsistent, and distributed medical knowledge be represented in CDSSs and managed in automated reasoning and decision-making?**

  In the medical field, knowledge is generated from different sources, such as different medical guidelines, clinicians' experiences, and patient data. This information is often uncertain, incomplete, inconsistent, and distributed. As a consequence, traditional logics that are monotonic, i.e., new information can not alter conclusions previously made, are not suitable for automated reasoning in a complex medical situation that requires reasoning over time due to the organization of care. As a consequence, methods that can handle non-monotonic reasoning, such as possibilistic logic, multi-agent systems (MAS), and formal argumentation, are more suitable for the medical domain. There is a need to develop methods that combine the advantages of these approaches and further improve on them for automated reasoning and decision-making in interaction with clinicians who are using CDSSs.

- **How can informal medical knowledge be translated into formal information that a computer can use, and how can the system architecture be modified so that the knowledge management and maintenance can be conducted through end-user development of CDSSs?**

  Despite the potential benefits CDSSs have, there exist a number of challenges hindering the implementation of CDSSs that rely on EBM knowledge such as: (1) the knowledge acquisition bottleneck, both in the initial formalization and knowledge engineering process, and also when research produces new medical knowledge; (2) the reusability of code for more efficient development of new CDSSs for different diseases; and (3) the ability to customize to the routines at different healthcare - providing organizations and different national medical guidelines. These bottlenecks occur because of the medical field's complexity and continual, rapid changes. Due to the hard coding, most of the existing CDSSs do not have enough adaptability to catch up with the constantly evolving medical field. Developing new CDSSs requires vast efforts from knowledge engineers and software engineers. Thus, there is a need to quantify and represent the informal knowledge accurately and, more importantly, to easily keep it updated with the changing situations. How to modify the system architecture so that the system can be developed directly and primarily by end-users rather than by knowledge and software engineers becomes a natural focus of research, since it is the medical domain experts who have the most updated knowledge.

- **How can learning and skill development be detected when a medical clinician interacts with a CDSS, and subsequently how can a CDSS offer tailored support so that the medical clinicians advance their knowledge and skills?**

  As mentioned earlier, one important purpose of CDSSs is to support learning and skill development, which is typically manifested in behavior. Person-tailored support for educational purposes is not often seen in CDSSs but could be highly valuable. For that, the users' knowledge and skill levels need to be adequately detected and followed over time to offer user-tailored support afterward. The third research question is thus how to accomplish these goals by developing methods that are based on knowledge about how novice and expert clinicians reason and make clinical decisions.

1.1 Objectives

The three research questions lead to three objectives:

Objective 1. From a computational perspective, to develop theories, methods, and algorithms based on possibilistic logic and argumentation framework in MAS for representing and managing medical knowledge pervaded with uncertain, incomplete, inconsistent, and distributed information and for automated reasoning and decision-making using CDSSs.

Objective 2. From a software engineering perspective, to modify the existing CDSS structure, develop methods and instruments to allow for the end-user development of the CDSSs, facilitating knowledge elicitation and transfer, and reuse CDSS modules with greater adaptability to work on breaking the bottlenecks.

Objective 3. From a human perspective, to explore the possibility of detecting reasoning patterns to measure users' skill levels and development, which could be used for providing personalized educational support through CDSSs.

In the pursuit of these objectives, it will be taken into account how clinicians learn, reason, and make clinical decisions for the purpose of providing tailored decision support in the medical field.

1.2 Theory and Methods

The research spans theory and implementation, including knowledge representation and reasoning with uncertain and incomplete information, software development, and qualitative end-user studies in clinical practice. The theoretical base includes formal theories from the field of non-monotonic logics for representing information and reasoning with the information, such as possibilistic logic and formal argumentation theory; literature studies on how humans conduct reasoning in clinical practice and analysis of data utilizing activity-theoretical models. Moreover, dialogue systems, MASs, and ontology technology have been explored.

1.3 Medical Context

The particular field of study in focus for this thesis is dementia. Dementia is a broad category of brain diseases causing a long-term and gradual decrease in the ability to think and remember that is severe enough to affect daily life. Typical dementia symptoms include emotional instability, language and other cognitive problems, and a decrease in motivation [20]. A dementia disease causes decreases in mental functioning and should be manifested in a greater decline than one would expect from aging [19]. Up to 70% of dementia cases are Alzheimer's disease [20]. The medical domain is very complex. In particular, the dementia domain is characterized by uncertain, incomplete, and inconsistent clinical information, in addition to its changing nature as a progressing and deadly disease. For example, Alzheimer's disease has mixed origins, and many phenomena remain unexplained.

In this thesis, for simplicity, we use the term "clinician" to refer to medical clinicians, e.g., physicians and nurse practitioners. They can all be the users of a CDSS. Clinicians are divided into experts and novices, according to their expertise level concerning their medical professional knowledge in a particular domain. Further, some medical domain experts have in this research been authorised to participate in modelling the knowledge using a content management system (CMS). They are also considered end-users, participating in an end-user development process.

1.4 Overview of the Contributions

Theories within the non-monotonic logic field are useful for dealing with the complexity of the medical field. In this work, possibilistic reasoning and argumentation methods are applied to deal with uncertain, incomplete, and inconsistent knowledge. Also, the MAS approach is applied to handle the distributed data to allow agents to make decisions collaboratively. Another reason for choosing MAS is to simulate the dialogue between expert and novice clinicians

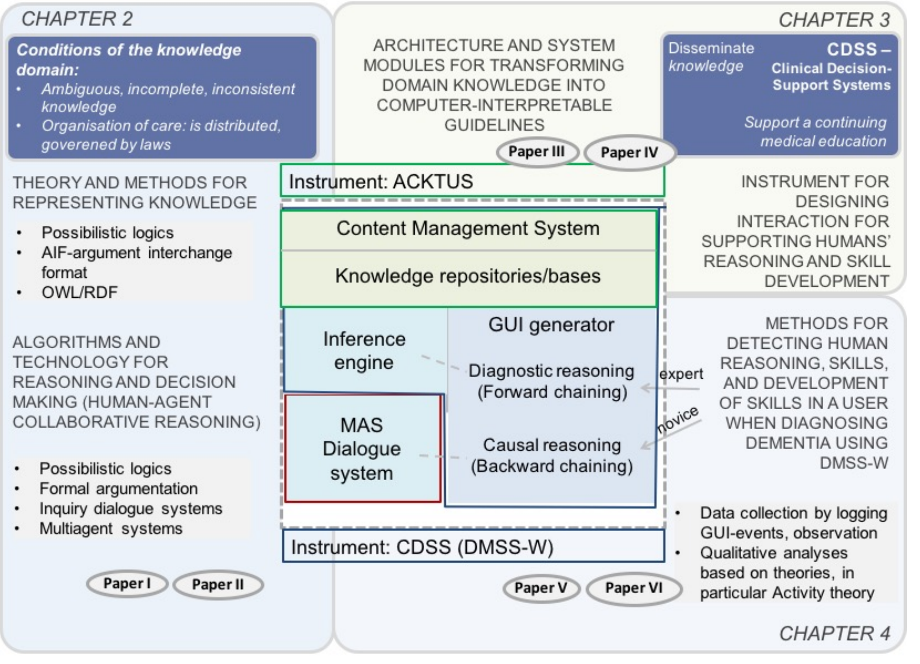

Figure 1: Summary of this thesis structure. Paper I - VI denote the six papers included in this thesis.

for collaborative reasoning. In this way, the above-mentioned data are properly handled in the implementation of the MAS (Paper I and Paper II).

This thesis digs deeper into the development of the Dementia Diagnosis and Management Support System - Web version (DMSS-W) and proposes a Semantic Web-based software structure for knowledge engineering, interaction design, and reuse of modules in new CDSSs. This Activity-Centered Knowledge and Interaction Tailored to Users (ACKTUS) system includes a CMS, a knowledge base (KB), a graphical user interface (GUI) generator, an inference engine, and a core ontology that is shared between those parts as well as different projects. Through ACKTUS CMS, the medical domain experts as users are given the opportunity to participate in and affect the development of the system. Consequently, the user's input can be directly incorporated to the user interface and also used by the inference engine without changing or recompiling any code or redeploying the system (i.e., "end-user development") (Paper III and Paper IV).

DMSS-W is applied in clinical practice for the purpose to investigate qualities of the implemented formal theories, methods and interactive reasoning support (Paper IV, V, and VI). A particular focus is set on how to support novice clinicians develop expert reasoning strategies and how to provide tailored support for different reasoning styles. For that focus, one first needs to determine the users' skill level. Cognitive science research offers an underlying theory about differences in reasoning strategies between expert and novice clinicians [90]. Based on this, the reasoning patterns are detected and used as the basis for providing tailored support to users (Paper V and Paper VI).

1.5 Organization of the Thesis

The following three chapters provide a background to the different complementary research domains in the introduction and underpin the results of this thesis (Figure 1). Chapter 2 corresponds to Objective 1 and provides a background of possibilistic logic, MAS, and argumentation, which are suitable for dealing with uncertain, incomplete, inconsistent, and distributed knowledge. Chapter 3 corresponds to Objective 2 and provides a background to knowledge engineering and designing interactive reasoning support. The existing bottlenecks in developing and managing CDSSs are described. The modified structure of ACKTUS is discussed in terms of facilitating the end-user development of CDSSs. Chapter 4 corresponds to Objective 3 and provides a background on the human perspectives of the knowledge, learning, and reasoning, as well as preliminary results on detecting reasoning patterns when clinicians are using DMSS-W. As such, it serves as a base for future research directions. Chapter 5 summarizes the contributions of this thesis and the included articles.

Chapter 2

Knowledge Representation and Collaborative Reasoning: the Computational Perspective

This chapter presents methods for managing uncertain, incomplete, inconsistent and distributed information, and for allowing the clinician to participate in a reasoning process in collaboration with the system. In particular, ontology and possibilistic logics as knowledge representation methods are introduced. For reasoning and decision-making, MAS and formal argumentation are introduced. The chapter provides a background for the theoretical and computational contributions of this thesis.

2.1 Knowledge Representation

Representing and managing knowledge is essential for clinical decision-making. For that, ontology is used in this research, from which the information is extracted for logic-based reasoning, e.g., using possibilistic logic. *Ontology* is a concept originally from philosophy and borrowed by computer science to define the types, properties, and interrelationships among the entities that exist in a particular domain. As an explicit specification of a shared conceptualization, an ontology is a knowledge model that represents a set of concepts and the relationships among these concepts within a domain [46]. It is one of the most successful ways to represent medical knowledge [27, 79, 110] because it helps capture medical knowledge in a formal but straightforward, powerful, and incremental manner, and it can be easily applied in the reasoning process of

CDSSs [97].

The use of formal ontologies is basis for the Semantic Web, which was first introduced by Berners-Lee and colleagues [15], to allow knowledge to be shared and reused on the internet across application, enterprise, and community boundaries. An example of a formal ontology is the Argument Interchange Format (AIF) [23] developed for the representation and exchange of arguments and their information between various applications. It contains a consensus "abstract model" agreed upon across the research fields of argumentation, artificial intelligence, and MAS. An extended version of AIF is integrated into the ACKTUS ontology.

The International Classification of Function, Disability, and Health (ICF)[1] developed by the World Health Organization is another example that provides concepts and their relationships, suitable to be represented as an ontology. In addition, the ICF defines its concepts, which is useful for clinicians in a continuing medical education process. ICF is used as the base for the *core ontology* (for details, see Chapter 3) that builds the KBs in the research presented in this thesis (Paper III), where the instances of the core ontology are inputted by the experts.

A number of programming languages have been defined that provide the basic machinery used to represent ontologies in the Semantic Web context. Resource Description Framework (RDF)[2] is a standard model for data interchange on the Web. It uses the triple format $\langle subject, predicate, object\rangle$, which is a standardized way of describing things and their relationships. The Web Ontology Language (OWL)[3] builds upon description logics [4, 5, 21] which is a knowledge representation language used for representing the domain knowledge in a structured and well-understood manner. It is based on concepts and roles, but differs from its predecessors, e.g., semantic networks and frames, in that it is equipped with formal, logic-based semantics.

OWL and RDF are built upon monotonic logic and can be used for reasoning. A typical example of monotonic logic is as follows. *A is a type of B. B is a type of C. Then A is a type of C.* Another example from the medical domain is: *A person is diagnosed with Alzheimer's disease. Alzheimer's disease is a type of dementia. Therefore, the person can also be diagnosed with dementia.* This is the typical setting of the *monotonic* inheritance networks [4]. However, the medical domain is more complex than this. For instance, since Alzheimer's disease is gradually progressing, a person can have Alzheimer's disease during an early stage when dementia cannot be diagnosed. Thus, a more powerful type of logic that can deal with *non-monotonic* reasoning is needed, which can manage a situation when new information is added leading to that old conclusions no longer hold.

As seen in Paper I, defeasible logic with integrated preference levels [16, 115]

[1] http://www.who.int/classifications/icf/en/
[2] https://www.w3.org/RDF/
[3] https://www.w3.org/OWL/?

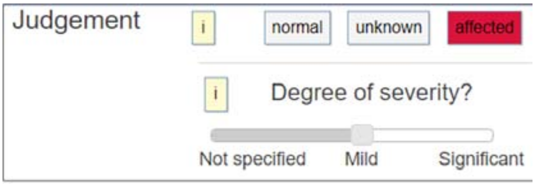

Figure 2: The symptom *judgement* and its status.

has been used for this purpose, but it could not adequately handle knowledge uncertainties [24]. Probabilistic logic, such as Bayesian inference, is known as a method to deal with uncertainties. However, the EBM data is often incomplete and does not provide a statistical base for applying statistical methods to support diagnosis in a particular patient case. Also,it is a quantitative method, whereas the judgments from clinicians are often qualitative, which cannot be easily translated into numerical values that represent the probability of each statement. For instance, in the medical field one often uses "It is possible that...", "I guess that..." during the decision making. Moreover, when utilizing the probability calculus, it cannot solve some problems [30], such as: (1) If the probabilities of p and $p \rightarrow q$ are both α, one can not conclude that the probability of q is also α. It means probabilistic reasoning cannot maintain probability bounds across inference steps. (2) When the probability of a rule is known, the probability of the counterpart of this rule cannot be derived directly.

In this thesis, therefore, we opt to use *possibilistic logic* [29] to capture uncertain information and degrees of confidence in knowledge sources. In this way, it can properly treat possibilistic uncertainty and incomplete knowledge. A key feature of possibilistic logic is that the strength of a conclusion is the strength of the weakest argument used in its proof. Possibilistic logic expressions are classical logic formulas associated with weights interpreted as lower bound of necessity degrees [30]. Weight p estimates to which extent it is certain that the possibilistic belief B is true: $N(B) \geq p$. In some cases, B can also be a goal, where p is then the priority level of goal B [62]. Here p is a number drawn from the range 0 to 1, where 0 means totally uncertain, and the closer to 1, the more likely is the situation that the belief is true, i.e., the confidence increases.

The possibilistic approach presented in Paper II evaluates arguments based on the reliability of the knowledge and information source. Reasoning with probabilities is not feasible here since the statistical distributions provided by evidence-based medicine (EBM) are not coherent. The data vary in their distributions, and there are overlapping conditions. In contrast, reasoning with possibilities does not directly rely on statistical information.

In this thesis, possibilities are extracted based on clinical guidelines, which are developed by domain experts by interpreting available evidence-based studies. As such, a possibilistic approach is more applicable since it follows how clinicians interpret EBM to make it applicable in individual patient cases. The following is a concrete example of how to map certain, incomplete, and uncertain data to possibilistic logic. In Figure 2, *judgement* is the patient's ability to decide which behaviors are appropriate under what circumstances. There are three choices for this symptom: [*normal*, *unknown*, *affected*]. For *affected*, there are successors to tell the severity of the level: [*not specified*, *mild*, *significant*]. *Normal* is a certain information, *unknown* is an incomplete information and *affected* with each of its successor is an uncertain information. From these data, the following possibilistic beliefs can be extracted:

- From *Judgement-normal*, $(\top \to \neg Judge, 1)$ is generated;
- From *Judgement-unknown*, both $(\top \to Judge, p)$ and $(\top \to \neg Judge, p)$ are generated, such that $(0 < p < 0.5)$;
- From *Judgement-affected*,
  (a) If *not specified* is chosen, both $(\top \to Judge, 1)$ and $(\top \to Judge\_n, p)$ are generated, such that $(0 < p < 0.5)$;
  (b) If *mild* is chosen, both $(\top \to Judge, 1)$ and $(\top \to Judge\_m, 1)$ are generated;
  (c) If *significant* is chosen, both $(\top \to Judge, 1)$ and $(\top \to Judge\_s, 1)$ are generated.

2.2 Multi-Agent Systems

Researchers from different disciplines (e.g., computer science, cognitive science, artificial intelligence, etc.) have produced various definitions of an intelligent software agent from different perspectives (e.g., [49, 73, 99, 106, 113, 114]). One of the most accepted definitions was proposed by Wooldridge [113]: "an agent is a computer system that is situated in some environment, and that is capable of autonomous action in this environment in order to meet its delegated objectives." MAS consists of a set of typically distributed intelligent software agents [114]. Briefly, MAS can be understood as a computer system that comprises two or more interacting agents within an environment.

In some domains, different people and organizations do not share the same goals and proprietary information. Each organization needs its own system to model its affairs because it has to reflect its priorities and capabilities. In such cases, a MAS is required. For example, medical data can be distributed in different places and used by different kinds of professionals involved in a medical process [9]. To make a decision, all the available data are needed. Each place only contains partial data for making the decision. However, the data can not be easily combined for different reasons. For instance, the data may be expressed or stored in different formats (e.g., 0/1 is used in one database, and Y/N is in another one; or one is in XML documents, and another is in a database), or simply because of the privacy or security reasons. Indeed, medical privacy is a concern in many countries, e.g., HIPAA is a strict American law pertaining to medical privacy. Many clinics, offices, and hospitals err on the side of caution when trying not to break any HIPAA rules or unnecessarily exposing protected patient health information. Therefore, MAS is used, which uses a set of agents. The agents represent different viewpoints, and they work together to solve a problem which cannot be answered individually. The way they work together is through a dialogue system, which is known as collaborative reasoning in this chapter. Since the original data is inconsistent, there could be some conflicting arguments generated during the dialogues, although the agents share the same goal (to decide whether the patient has the disease or not). The necessity to use MAS is analogical to the situation where two persons jointly make a decision in daily life. Since it is impossible for one person to automatically know the thoughts stored in another person's brain, they must talk to communicate. Similarly, one agent does not have access to the KB of other agents, and thus dialogues within a MAS are necessary.

Even in situations where the data are not physically distributed, a MAS can be useful [106]. The most obvious benefit is the accelerated operational speed by the parallel computation enabled by the multiple agents. Meanwhile, the system becomes significantly more robust due to the additional agents that participate in the computation. The scalability of a MAS is also a benefit. As the MAS is intrinsically modular, it is much easier to add new agents in case of necessity, as compared to a monolithic system. It, in turn, leads to the simpler programming of the MAS.

MASs have been used in the medical field, e.g., medical data management [108], patient scheduling [74], remote healthcare for older adults [26, 43, 61, 107], and CDSSs [44, 104]. In fact, MASs have been found to be highly suitable for the healthcare domain [22] which is a vast field where decision-making is frequently made in a distributed way, requiring communication between different departments, clinicians, and patients. A MAS makes it easier to make decisions and take actions, which can greatly improve healthcare [22].

The motivations for using a MAS in this research are twofold: (1) The data are physically distributed in different places and cannot be combined into one common KB for reasons such as privacy. However, all data are needed to reach a decision. (2) It may provide novice clinicians with a tool for diagnostic reasoning that allows them to begin in a tentative hypothesis, mirroring the causal characteristics of medical reasoning in novices described in Chapter 4, but preventing essential aspects from being missed, which is a risk with causal reasoning. As such, the reasoning mechanism of MAS complements the inference engine that applies primarily forward-chaining reasoning, presented in Paper IV, and may contribute to educating the user by leading the multi-agent dialogue into areas that the clinician otherwise may miss. The system educates the users by providing them transparent explanations of the reasoning and decision making used to diagnose patients.

The agents have the social ability to communicate and, hence, generate dialogues. Different types of dialogues were defined by Walton and Krabbe [112]. These include *information-seeking dialogues* (aiming to transfer knowledge and information from one party to another), *inquiry dialogues* (generating new knowledge as conclusions and decisions), *deliberation dialogues* (generating decisions on how to act), and *persuasive dialogues* (aiming to resolve conflicting viewpoints through compromise). Dialogues follow *dialogue protocols* that set the conditions for a dialogue [3, 77], e.g., the number of allowed participating agents, how the order of initiative is organized, and conditions for when dialogue can be considered closed and completed.

In this thesis, inquiry dialogues are in focus as a means to collaboratively generate new knowledge in a patient case. The goal of an inquiry dialogue is to prove or contradict and possibly falsify the hypothesis in a proof process of a collaborative reasoning, and leads to decision-making. A starting point is that the participating agents possess different but potentially complementary knowledge that may contribute to a reasoning process differently. This means an agent can also learn from the other agents while participating in a dialogue. Earlier research on MAS that apply inquiry dialogues is very sparse [16, 17], and rarely deployed in practice [52].

In a similar way to [16], we use two kinds of inquiry dialogues in the framework: *warrant inquiry* (wi) dialogue and *argument inquiry* (ai) dialogue. Ai dialogue generates the argument that can be used in a wi dialogue (i.e., if all the premises of a rule can be proven to be true, then an argument is generated). Wi dialogue contains 0 to n ai dialogues. Wi dialogue generates new knowledge by comparing these arguments. In the further developed solution presented and implemented in this thesis, these two types of dialogues are nested within each other, as schematically shown in Figure 3.

A dialogue between participating agents is actually a sequence of *moves*. The move concept was first introduced in [84] and is similar to the concept of a *speech act* in literature [76]. The moves are used for communicating to each other; a dialogue allows for different types of moves. Typical examples are *open*, which is for opening a dialogue, *assert* for introducing new knowledge, and *close* for closing the dialogue (see [16], for instance).

In agent dialogue frameworks, protocols are usually presented to determine the agents' next moves [16, 58, 75, 95]. In this thesis, a protocol is regarded as a function that, given a particular type of dialogue, a specific move is returned according to its belief base and previous moves they have already made.

2.3 Formal Argumentation

In this section, we describe how the arguments are generated and evaluated via dialogues between agents in the MAS. Argumentation is an interdisciplinary research topic, engaging researchers primarily in the fields of artificial intelligence, law, linguistics, and philosophy. The motivation of applying argumentation is

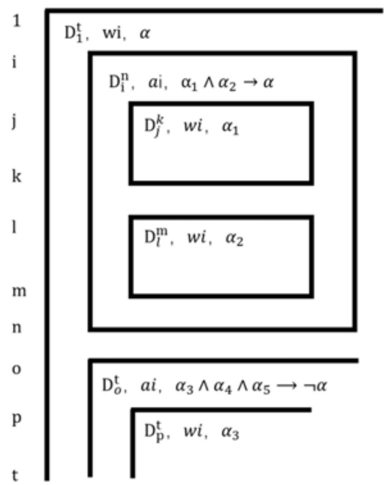

Figure 3: Schematic diagram of the nested inquiry dialogues. The text in the left column represents the timepoints of the moves.

to deal with inconsistent information [13]. The generated arguments are evaluated to reach a conclusion, i.e., to find a winning argument in a dispute. This evaluation can be completed by following different argumentation semantics [31].

An argument typically consists of two items: the tentative conclusion (the *claim*) and a set of premises that support the claim [60]. The premises can contain potentially defeasible facts and rules. Dung developed in his seminal work [31] an abstract argumentation framework (AAF) where the central notion of the acceptability of arguments is captured in a general way. Concretely, in an AAF, there is a directed graph such that the nodes are the arguments and the arrows represent the attack relation. Based on the graph, Dung defined different semantics for the acceptability of arguments. There are extensive works on extending AAF, such as value-based argumentation framework [12], bipolar argumentation framework [2] and preference-based argumentation framework [1]. These frameworks include the *evaluation* of arguments by calculating the acceptability of arguments.

In our CDSS, since the data are distributed in different locations, a MAS is employed. During the reasoning, to reach a joint decision (which is known as an argument in our approach), the two agents communicate and generate dialogues. In the dialogue, potential arguments are brought up and verified potential arguments become real arguments. Then, backward reasoning is used, which starts with a conclusion, finds rules that support or attack the conclusion, and checks if all the premises of the rules are fulfilled. If yes, a real argument is generated. In case several arguments are generated regarding one conclusion, the following algorithm is used to determine which wins. First, the arguments are divided into two sets: supporting or attacking the conclusion. The set with the largest possibilistic value wins. If the largest possibilistic values of the two sets are equal, we count how many arguments are contained within each large-value set, and the one with more arguments wins. If the counts are still equal, the second largest possibilistic values are compared in terms of the values and counts. The evaluation continues until a winner is found, or the competition ends up with a draw (result unknown). We note that the backward reasoning restricts the search for arguments to those whose individual inference steps are all relevant to the conclusion, compared to the forward reasoning [32].

The procedure to generate arguments via dialogues can be viewed as an AND-OR-tree. There are two kinds of nodes: a wi node whose element is a literal (an atom α or its negation $\neg\alpha$) and an ai node whose element is a rule, corresponding to the wi and ai dialogues, respectively. The root node is a wi node. Its element is a hypothesis which is a literal, from which the rules whose conclusions are this hypothesis or its opposite can be found. As a branch of the wi node, each rule generates an ai node. The relation between these ai nodes is OR. From each ai node, there are again wi nodes where each premise is a branch. The relation between these wi nodes is AND, meaning an argument is generated only when all the premises are fulfilled. Leaf nodes are always wi nodes. Each ai node (together with branches) can generate an argument if all the branches are proven true. Under a wi node, there can be several arguments generated, supporting or attacking the literal. The algorithm to decide the acceptability of arguments was described earlier. The backward reasoning can be seen as generating these trees in a top-down fashion, i.e., from root to leaf.

In a way, our approach is somewhat similar to Dung's work concerning the use of backward reasoning [32], but there are some apparent differences. (1) Dung's approach uses assumption-based argumentation, and only assumptions can be attached. However, in our case, the argumentation is based on possibilistic logic where all the facts and rules are defeasible. An argument can attack another argument by contradicting its conclusion. (2) The acceptability of arguments is different. In Dung's work, if the proponent can defend the argument by counter-attacking the opponent's attack, then it is acceptable. Our method takes more caution in this regard because we not only check if one argument can attack another argument using proper rules but also take a detailed look into the possibilistic values and their counts of the two sets.

Chapter 3

Knowledge Engineering and End-User Development: the Software Engineering Perspective

In the development of a new CDSS, informal knowledge needs to be translated into formal, computer-interpretable representation formats. Also, the knowledge in the CDSS needs to be dynamically maintained (knowledge maintenance). This process is resource demanding, both in the initial stages and while updating the contents as new knowledge appears [18]. The *knowledge acquisition* bottleneck is a well-known limitation of traditional, rule-based CDSSs [11, 41, 59]. In medical research, new knowledge is continuously created. Research shows that "there is stark evidence of a 13-17-year gap between research and practice in clinical care" [14]. This evidence indicates that effective methods for transforming new scientific results into clinical practice are lacking.

Another related challenge is *code reusability* when developing for a different disease [18]. Usually, it takes a lot of time and effort to develop a CDSS for a specific disease [8, 48]. If the code of an existing and well established CDSS can be reused and quickly developed into a new CDSS, it will save time and, potentially, lives. Especially, when a new disease suddenly breaks out, it will be beneficial to rapidly develop a CDSS accordingly. However, today most CDSSs are hard-coded and are not highly reusable.

There is also a difficulty when *customization* is needed to comply with the routines at different healthcare providing organizations, e.g., following different versions of national treatment protocols [10, 80]. Clinicians from different countries or with different background do not necessarily apply the same methods for physical examination or the same diagnostic criteria [92]. There may be considerably large variations in the way they interpret the data. An earlier study shows that differences between countries and organizations were observed in versions of validated assessment instruments and preferences regarding which international diagnostic criteria would be applied [66]. For example, some checklists are used in Sweden, but not in Japan. It is thus important to tailor a system to different countries or organizations and integrate it into their digital infrastructure in order to be widely used. However, the existing CDSSs are mostly designed for specific organizations and their particular requirements relating to their internal digital infrastructure.

This chapter first introduces some software tools developed for knowledge engineering and interaction design purposes. Then, the development of the user interface is described, where ontologies play a role. The last section demonstrates that the previously described challenges can be properly dealt with by involving authorized end users (domain experts) in the system development, minimizing the need for knowledge and software engineers.

3.1 Engineering and Maintaining Computer-Interpretable Clinical Guidelines

The transformation of medical knowledge into computer-executable formats is usually done by a knowledge engineer in collaboration with a domain expert who provides interpretations of the informal domain knowledge. Some tools are available for that, e.g., Protege [42] and various software for developing business intelligence. Different research groups have developed task-network modeling (TNM) languages for representing clinical practice guidelines, treatment protocols, workflows, and their decision-making tasks [36, 78, 85]. In order to have computer-interpretable representations of the guidelines, Asbru [101], EON [109], GLIF [93], GUIDE [96], PRODIGY [53], and PROforma [37] were developed. To date, however, there are not yet standards for these approaches in terms of goals, expression language, data interpretation, guideline models, or other requirements [94]. First, their intended scopes were very different. For instance, Asbru and PROforma have deliberately not included methods for representing static knowledge such as medical concept models. That is because they were developed by different research groups with different interests. They also used various guideline models. PRODIGY, EON, and GLIF used a scenario-based approach, where a guideline was organized as a collection of clinical contexts, and the users could select contexts from relevant clinical actions. Indeed, clinical scenarios are intuitive, which may be easier for domain experts to manage [53]. Finally, how to represent and utilize the clinical goals varied a lot among these systems. Asbru, EON, PROforma, and GUIDE represented goals formally and allow reasoning, but GLIF and PRODIGY did not.

However, the methods presented can only generate computer-interpretable clinical guidelines, which are fragments of CDSSs. They are important in the

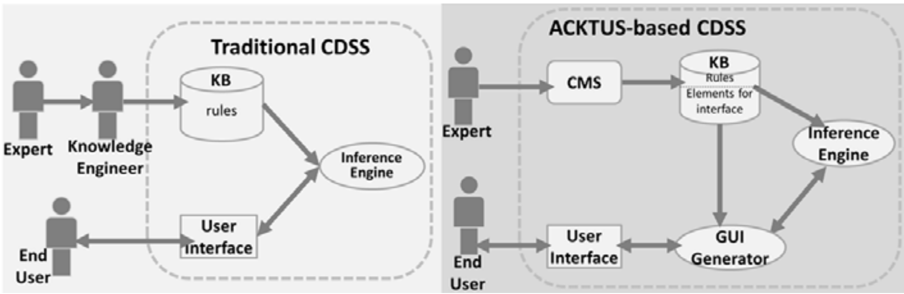

Figure 4: Traditional CDSS architecture and the extended ACKTUS-based architecture presented in Paper IV.

applications; however, our aim is beyond that. We would like to systematically develop methods suitable for generating entire CDSSs. Demonstrated below is our approach, which modifies the traditional CDSS architecture. A rule-based CDSS typically contains three modules, as exemplified in Figure 4) [45]. The three modules in the left panel are:

- A *knowledge base (KB)*, which stores rules. It is a centralized repository or a database of related information about a particular domain;

- A *user interface*, which allows the system to display the decision results to the user and collects user input;

- An *inference engine*, which applies the rules stored in the KB to patient data retrieved from the user interface or electronic health record system to generate patient-specific recommendations.

ACKTUS is used in the research presented in this thesis for knowledge engineering and maintenance. It makes use of an ontology-based approach to model medical knowledge. It represents clinical information formally and allows reasoning, which finally leads to diagnostic decision-making. An important feature of our ACKTUS-based approach that differs from other methods is that it facilitates end-user development to deal with the challenges of CDSSs described in the beginning of this chapter.

The ACKTUS architecture presented in this thesis is shown in the right panel of Figure 4 and is designed with the following two additional modules:

- A *content management system (CMS)* attached to the KB, which is built on Semantic Web technology to achieve modularity, reusability, customization, and the possibility to allow medical experts to model the knowledge as well as structuring the information that builds up the design of the user interface;

- A *GUI generator* that automatically generates the user interface whenever the user logs in so that the interface is synchronized with any updates in the KB without a software developer's intervention.

3.2 Using Ontologies for Developing User Interfaces

Ontologies can be used for representing a knowledge domain. In philosophy, the ambition is to develop a "true" representation of entities and events in the real world, e.g., diseases and body parts. However, in practical applications of ontologies, sometimes there is a need to represent even something that does not exist, for instance, an amputated leg, which is useful information in the medical domain. This usage is quite different from the original meaning of ontologies in philosophy. Software can also be defined using ontologies, where a common ontology helps the different parts of the system work harmonically. Another application for ontology is to characterize users, their roles, and their preferences when they use the system.

In this thesis, ontologies are not only very important for the above-mentioned aspects but also useful for enhancing the user interface of CDSSs. Ontologies can be used to either provide a single functionality in a user interface or to construct the whole interface. In literature, there are different types of ontologies used for these purposes [91]. Regarding how the ontologies can enhance the interface, there are three categories: (1) improve the interface visualization; (2) improve the user interaction; and (3) improve the development process [91]. The system presented in Paper IV mainly belongs to the third category, but also overlaps with the other two.

There are systems using ontologies to identify and track requirements from the real world to generate the interface [38, 105]. Some systems combined the system ontology with user ontology to create personalized user interfaces, such as OntoWeaver [65]. The development of user interface consumes 50% of the total efforts [81]. Thus, some systems used ontology for reusing user interface components [47]. Ontology technology was also commercially used to generate interfaces[1].

Based on [50, 91], most systems employed static ontologies, meaning the ontologies could not be modified or extended. In other words, in Semantic Web applications, the interfaces cannot be altered by the users. By contrast, the method presented in Paper IV enables the users to edit the interface by modifying the structure and content using the CMS and, subsequently, to synchronize the interface with these changes through the GUI generator. The reason this can be achieved lies in the fact that the domain ontology included in the system is editable. To some extent, it fills the gap currently existing in other reported works, since it meets the need for a more intuitive interface for

[1]See, for example, http://mitosystems.com/

revising and manipulating ontologies [39]. Through the GUI generator and the editable ontology, the user interface of the CDSS can be developed by the end users.

3.3 End-User Development of CDSSs

It has been shown that when medical experts are given control over the modeling task, they tend to be more strict in their modeling and discover ambiguities that they need to manage [72]. Moreover, when they are given tools for rapid prototyping, they can design, according to their vision of a knowledge-based support system, the content and interaction with the application, which they can also evaluate by involving colleagues or patients [68, 70, 82].

The TNM systems are designed for facilitating domain experts' participation in the development, in particular, to visualize clinical workflows. However, the typical knowledge acquisition procedure is to engage medical experts working together with knowledge engineers in their modeling of the content, while the software tools for knowledge engineering are still designed for the engineer as the user rather than the medical expert. Some experiments with intermediate modeling steps have been conducted to involve medical experts in more hands-on modeling [6, 7, 33, 100, 102]. Still, it was found that extensive time is required from the medical experts. It has been shown that the method and the interface for non-programmers to use for the modeling of rule-like structures, need to be very simple and intuitive [6, 7]. There are also some attempts to allow end users to participate in the interface design [54], but they are limited to designing the interface without considering the knowledge engineering underlying the CDSSs.

The research presented in this thesis aims to allow authorized end users, i.e., domain experts (denoted "Expert" in Figure 4), to be involved in the system development, as they know the newest developments in medical knowledge. Compared to how knowledge management is typically done, where knowledge engineers and software developers dominate the system development, the knowledge engineering is performed so that the technical bottlenecks depicted in Paper IV can be resolved.

In the ACKTUS architecture, the domain experts directly model the domain knowledge using the CMS, reducing the need for knowledge engineers. The knowledge is stored in the KB, and the GUI generator fetches data from the KB and generates the user interface automatically. Authorized domain experts are allowed to manage the content using the CMS and in this way affect the CDSS. The end users (including both authorized experts and other clinicians) assess the patient and fill in the patient symptoms by using the generated user interface. The inference engine module uses the symptoms obtained from the interface and the rules extracted from the KB to conduct reasoning.

Because the domain experts can directly model knowledge using CMS, the new knowledge can be included in the KB quickly and easily. Since the GUI

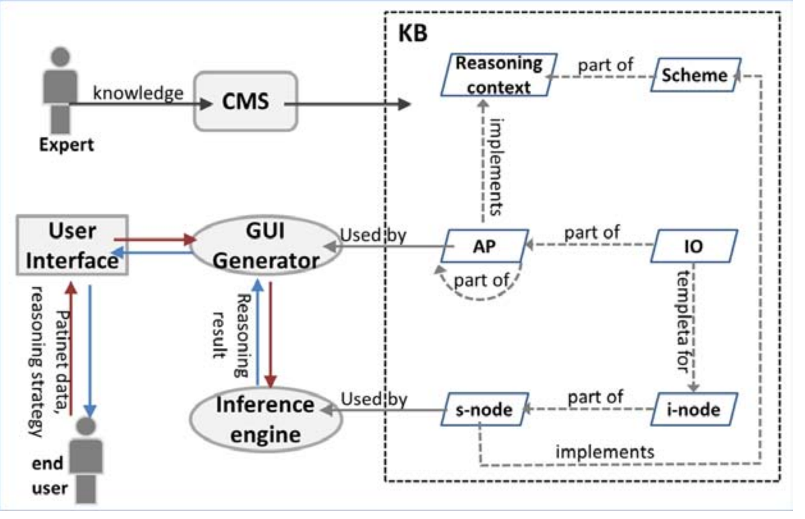

Figure 5: The workflow of knowledge engineering and automated processes that result in a CDSS that can be immediately verified by an expert and used by the end user.

generator automatically generates the user interface whenever it is loaded, the web page is synchronized with updates of the KB. As each part is developed separately, they can be reused individually in other CDSSs, and this feature is especially crucial when new diseases need to be managed. When the GUI generator generates the interface, it will check who the user is and only show user-related information. The user interface is tailored to each particular user, based on the user's organization and professional skill.

3.3.1 Knowledge Engineering and Interaction Design Using ACKTUS

In this subsection, we further elaborate on the ACKTUS-based system in Figure 5, illustrating how experts can represent knowledge and design the user interface and interaction. The domain experts with mark-up training can model knowledge with correct syntax and semantics. The knowledge can then be understood, interpreted, and utilized by the system through the *core ontology*, which defines some key classes that function as a universal data structure and shared vocabulary between different parts. It is extended with sub-classes and instances. The core ontology consists of three major parts with their corresponding functions:

- Patient information: a *concept-node system* consisting of a combination of the ICF and other medical terminologies, and extended with *scales* for evaluating the findings;

- Clinical knowledge: an extended version of AIF developed for exchanging arguments on the Internet [23], mainly consisting of the *scheme*, *information nodes (i-nodes)* and *scheme nodes (s-nodes)* from which rules are extracted;

- Interaction and GUI design: an ontology for GUI objects and their relations, mainly consisting of templates for information collection and structuring (*interaction object (IO)* and *assessment protocol (AP)*) and also for reasoning guidance (*reasoning-context* and *critical-question (CQ)*).

As shown in Figure 5, the instances of APs, IOs, and schemes are created by the domain experts through CMS. An IO with each of its scale values automatically forms an i-node. The i-nodes are combined into an s-node which is an instantiation of a scheme. A scheme is part of a reasoning context. The logic relations used to link these elements together are obtained from the CMS, which is, again, the experts' input. The classes of the core ontology have different properties that describe the classes in as much detail as necessary. For details, see Paper IV.

The GUI generator is a program developed using Java, jquery[2], and CSS. It searches the tree-structured data and extracts the data using a filter based on the properties *has-organization* of the data, the user's professional skill, and the selected language, until finally turns the filtered data into the actual interface. Importantly, the user interface is changed simultaneously with the contents in the KB without redeploying the website. Hence it is easy to extend its content.

When a user logs in, the GUI generator fetches the AP defining the application and retrieves all the sub-APs and -IOs contained in its hierarchy. The APs are used for generating the tabs and menus while the IOs are primarily presented as checklists or diseases. The values of a scale related to an IO are the assessment options of the checklist or the diagnosis options of the disease.

To realize the adaptability in the generation of user interfaces and results of the inference engine, four *"key AP instances"* and their relationships are defined and stored in the core ontology (exemplified in Paper IV):

> The **application AP** (in DMSS-W named "DMSS - Dementia Diagnosis and Management Support System"):
>
> - The **data capture AP** (in DMSS-W named "Data Capture"),
>   - The **reasoning-context-based AP** (in DMSS-W named "What to do next?");
> - The **diagnosis and intervention AP** (in DMSS-W named "Diagnosis and Intervention").

These key instances have dedicated purposes and fixed IDs in the application so that the GUI generator understands where to extract the relevant data.

[2]https://jquery.com/

However, the name, description and sub-AP and -IO of the key instances are modifiable using the CMS. The key instances include the instance of the top level AP (the root node in the tree structure) that defines the application. From the application AP, the GUI generator retrieves all the included APs and IOs that are the children and grandchildren of this instance. The application AP has at least two sub-APs: the *data capture AP* that dedicates for capturing the patient-specific data and the *diagnosis and intervention AP* for showing the diagnosis and intervention results.

The IDs of the four key APs are also stored in the corresponding interface so that the inference engine knows where to fetch and feed the data. The inference engine can attain *facts* from the *data capture tab* in the user interface to reason. An overview of the reasoning results is shown to the user in the *diagnosis and intervention tab*, where the user can explore the arguments for or against different diagnoses based on the different knowledge sources. The final decision is the end user's responsibility.

By defining the key instances, the GUI generator and inference engine can be applied in a new ACKTUS-based application without modifying the code. The bottlenecks depicted earlier have been resolved or relieved using the architecture in Figure 5, as described in Paper IV.

Chapter 4

Knowledge, Learning, and Reasoning: the Human Perspective

For CDSSs to offer user-tailored education, it is critical to understand how clinicians reason. A literature study has been performed focusing on how reasoning is typically conducted by novice and expert clinicians.

The results are used as the basis for designing the interaction with a CDSS (the collaborative reasoning and decision making), the formal theories underpinning the interaction designs, the inference engine, and the methods for detecting user skill development. The methods are applied in pilot case studies aiming for providing user-tailored educational support.

4.1 Learning, Reasoning, and Decision Making in Clinical Practice by Novices and Experts

Knowledge is the familiarity and understanding of things and processes, e.g., facts and their relationships, human beings, and procedures. It is obtained via experience, and more formally, by education. The knowledge acquisition of humans involves several cognitive processes: perception, communication, learning, and reasoning [28]. Learning and reasoning are two fundamental cognitive functions that humans possess. Learning is a set of processes for acquiring new knowledge and skills, which is a consequence of interactions between persons and their environment. According to studies in educational psychology, neuropsychology, and pedagogy, there are many types of learning. For example, children can learn through playing. In [51], learning is seen as a productive construction, namely the development of new neuronal connections in the brain.

Reasoning is to work through problems so that one can explain either why some conclusions have been drawn or what conclusions will be drawn.

In earlier research, propositional analysis techniques were used to examine the protocols of cardiologists in their diagnosis of acute bacterial endocarditis [86]. It was found that those physicians that had made more accurate diagnoses used pure forward reasoning through a network of causal rules. They started with the symptoms and reasoned until a diagnosis was reached. In contrast, those who made inaccurate diagnoses began with a high-level hypothesis and tried to verify the hypothesis using the symptoms; i.e., they applied only the causal rules that were relevant to the hypothesis.

From studies on how clinicians reason and make decisions [35, 86, 89], the following scenario can be derived. An expert usually first assesses the situation of the patients and builds up evidence for syndromes, before the reasoning about the diagnostic conclusion and the corresponding causes. In other words, the clinician collects information, e.g., symptoms from the patients, and subsequently makes a decision based on the clinician's rich knowledge and experiences [63, 86]. This *forward-chaining reasoning* is the typical expert reasoning pattern [86, 87, 90]. In contrast, novice physicians tend to apply causal reasoning already at the initial stage of the assessment process, beginning with the potential explanation of the situation. That is, the novice first makes a hypothetical diagnosis and then investigates features, which may or may not support the hypothesis. The physician may make necessary adjustments to the hypothesis at a later stage. This *backward-chaining reasoning* pattern is frequently seen among inexperienced clinicians who lack the ability to analyze a range of diverse symptoms critically, and have to rely on a hypothesis first and subsequently modify it based on the collected symptoms [87, 90]. The risk associated with this behavior is to miss some important information (e.g., the patient symptoms and diseases may be untypical). To summarize, experts usually conduct forward-chaining reasoning, whereas novices typically use backward-chaining reasoning [90].

In cognitive psychology, human problem-solving methods are well studied in several dimensions: forward vs. backward; knowledge vs. goal; data vs. hypothesis [88]. Forward-chaining reasoning refers to the situation that the problem solver work from the given information, i.e., from the data to a hypothesis or goal; whereas backward-chaining reasoning begins in hypotheses and leads back to the given information [86]. This practice is depicted in Figure 6, where the forward-chaining and backward-chaining reasoning are associated with expert and novice behaviors, respectively.

In practice, nevertheless, variations of behaviors are also seen. Experts can also exhibit some causal reasoning pattern, especially when they are conducting diagnosis while explaining to a medical candidate [35]. This variation notwithstanding, the reasoning pattern can still be used as a criterion to distinguish between experts and novices (see Paper V and VI).

The way medical professionals conduct clinical reasoning and decision mak-

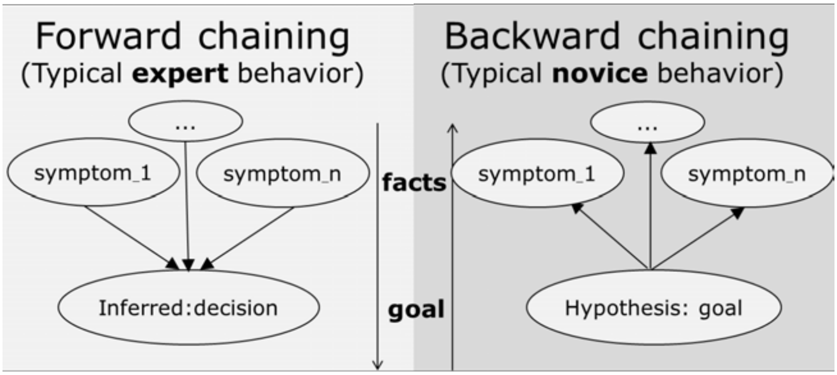

Figure 6: The comparison of forward reasoning and backward reasoning.

ing forms the foundation of the design and development of the DMSS-W. It fully supports two complementary reasoning processes [86]. When a clinician applies the *forward-chaining* reasoning method, clinical assessment and investigations are typically conducted before potential hypotheses are generated and evaluated. The corresponding procedure when using DMSS-W is entering all available information and findings in checklist format generated in the *data capture tab*, and then apply automated reasoning by activating the *diagnosis and intervention tab*. Based on the information available in the *data capture tab*, the system generates hypothetical diagnoses and their strengths, based on a set of international medical diagnostic guidelines. The results are presented to the user through the *diagnosis and intervention tab* as an overview of diagnostic conclusions and their strengths and support. If the information is not sufficient for a diagnosis, that is also presented, along with what information is missing. The tentative diagnoses may be conflicting, in case different diagnostic guidelines provide contradictory results. This information is also vital knowledge for the user who can then make a decision by selecting which medical source is preferred.

An alternative method that the user can apply to be guided in a *forward-chaining* procedure, if not familiar with the medical domain, is to use the *reasoning-context-based guide* called "What to do next?" in DMSS-W, which guides the user one step at a time towards diagnosis and intervention (Figure 7). When the user activates this functionality, the system generates checklists with a subset of symptoms to fill in. At each step, information about how to proceed and the sub-conclusions that can be made about diagnoses are provided. The sub-conclusions are answers to the critical questions that define each step. Finally, a list of supported hypothetical diagnoses is presented to the user for review.

The opposite strategy, typically seen in novice clinicians, is the *backward-chaining* causal reasoning which begins with a hypothetical diagnosis, e.g.,

Figure 7: Overview of the three steps of assessment following the reasoning context-based guide.

Alzheimer's disease. The risk with jumping directly to conclusions is to miss less common diseases, and, therefore, the interaction design of the CDSS promotes the forward-chaining reasoning strategy. The CDSS design allows the user to use the inference engine without conducting a thorough assessment. Then the user will be provided the overview of weak or unknown support for different potential diagnoses, with information about what patient information is missing for each potential diagnosis.

An alternative *backward-chaining* approach is a dialogue system starting with a hypothetical diagnosis, implemented using a MAS presented in Paper I and II. The user selects a patient (whose symptoms are stored in a client-side database) and a hypothesis about which disease the patient might have. The MAS is triggered and starts with the hypothesis to check the fulfillment of any rules that are supporting or contradicting the topic. During the reasoning process, alternative diagnoses are also taken into account.

4.2 Evaluating Learning and Skill Development

Activity theory was developed by the Russian psychologist Vygotsky and provides models of human purposeful activity [111]. It constitutes one of the most important theoretical foundations for human-computer interaction research. It describes human activity, including learning as social in its origin, and that activity is always mediated by tools. It fits the human-computer interaction perspective since the computer and its software are designed to function as tools in activities. Moreover, a computer may function as the tool to conduct social activities. Although we are alone in front of a computer, we are engaged in a collective activity, even when the activity is distributed in space and time [98]. As a consequence, a computer-based tool such as a CDSS can be used for computer-supported collaborative work and learning, which explores the social aspects of learning through work-based learning [25]. Computer systems are found to support human learning [55, 56], where learning takes place as part of the interaction with the computer. Activity theory emphasizes that humans develop by being active in purposeful activities.

An early study of DMSS showed differences in how clinicians complied with the suggestions provided by the CDSS [67]. A qualitative analysis was done, which indicated that lack of knowledge about dementia diseases could be the reason for non-compliance. It was interesting to investigate further whether individual clinicians' reasoning and skill development could be detected through patterns discovered when they use DMSS-W. If this could be done, then one can develop personally-tailored support for continuing medical education. In this thesis, the following research questions are addressed (some partially answered in Papers V and VI): (1) Can different reasoning patterns be detected based on log data collected when DMSS-W is used? (2) Can learning and skill development be detected? (3) Can novices and experts be distinguished by how they use the system? (4) Can it motivate the development of personalized reasoning support? (5) How can the support be tailored to these different categories of users?

Three complementary studies were conducted for the purpose of (1) developing methods for detecting reasoning patterns and skill levels (Paper V), (2) for detecting potential development of skills and knowledge (Paper VI), and (3) for identifying type of support that the clinician may need, relating to the skill level of the individual clinician [69].

DMSS-W is designed based on activity theory. A main assessment protocol contains an ordered list of nested sub-protocols, which builds the menu system in the user interface. In practice, the protocols represent a hierarchy of sub-actions in the assessment task with particular goals (see Figure 8). As a consequence, the click events are interpreted as follows (Paper V):

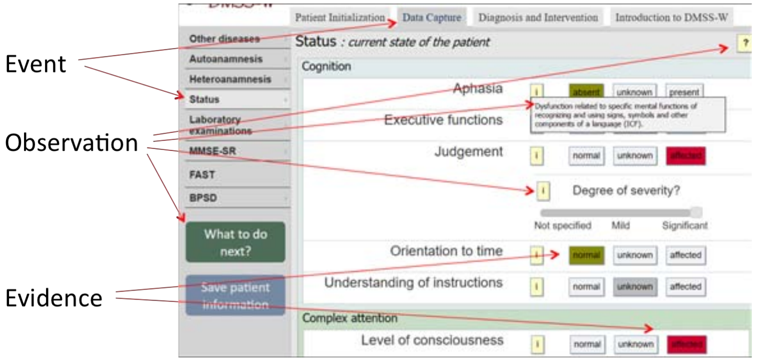

Figure 8: Illustration of "event," "observation," and "evidence."

- Category 1: When a user clicks on a menu or a tab (corresponding to a protocol), it is interpreted as the user shifts focus to the topic of the selected protocol. Specifically, a shift from *diagnosis and intervention* to *data capture* with subsequent logging of new evidence representing increased awareness of or new insights about missing information.

- Category 2: A click on a value associated with a phenomenon or symptom generates information corresponding to a decision made by the user, which is represented as evidence associated with the ongoing event.

- Category 3: A click on a button which leads to the help function is logged as an observation associated with the event. The help can be either about how to use the system or explanations and definitions of concepts related to symptoms and diseases. There is a guide button *"What to do next"*, which is a special help button. A click on it represents either a lack of knowledge about what to do next or an intention to speed up the process.

Each event is logged with information about time, actor, type of activity, and evidence. A semi-automatic analysis of the log file leads to a categorization into three types of sessions: (1) the patient case is complete; (2) partial information is entered about a patient case; and (3) no patient-specific information is entered, meaning the user navigates around the user interface without assessing a patient. Log data was collected from 29 physicians, where selected users were further analyzed.

Two case studies were conducted. The first was to apply and evaluate the methodology to users who had categorized themselves as either novice or experts in dementia, and the second aimed to follow users of the two categories over a number of patient cases to qualitatively and semi-automatically discover behavior patterns and potential changes of behavior (See Paper V and VI).

Furthermore, a third case study was conducted, using information and observations from four physicians in training to become experts in healthcare [69].

Act. level	Activity description	Level of independence			
Activity	**Dementia Assessment and Diagnosis**	**Person A**	**Person B**	**Person C**	**Person D**
Level 0	Understands the purpose of the activity but does not participate	A	A	A	A
Level 1	Contributes with partial information	A	A	A	A
Level 2	Contributes with domains of information	ZPD-S	ZPD-S	A	A
Level 3	Completes the assessment	ZPD-S	ZPD-S	A	A
Action	**Patient Information Management**				
Level 0	Uses the system without storing patient information	A	A	A	A
Level 1	Manages to create a new patient case and store this	A	A	A	A
Level 2	Manages to store existing case and load settings for a new case	A	ZPD-S	A	A
Level 3	Shows full understanding of the patient information management	ZPD-S	ZPD-S	ZPD-S	A
Action	**Data Capture Related to Dementia Knowledge**				
Level 0	Explores each symptom related to dementia assessment as having a goal in itself, i.e., executed as activities	A	A	A	A
Level 1	Assesses symptoms having partly been integrated in activity as actions among other actions	A	A	A	A
Level 2	Assesses categories of symptoms and view the different sources in a holistic sense	ZPD-S	ZPD-S	A	A
Level 3	Fully familiar with the different symptoms and their role for dementia assessment	ZPD-S	ZPD-S	ZPD-S	ZPD-S
Action	**Data Capture Related to Using the Instrument DMSS-W**				
Level 0	Able to use the information objects, their definitions and their values	A	A	A	A
Level 1	Understands the more complex objects and their severity scales, is able to use the GUI as guide for data capture	A	ZPD-S	A	A
Level 2	Understands the difference between necessary and optional information	ZPD-S	ZPD-S	ZPD-S	A
Level 3	Fully familiar with the data capture objects in the GUI and their role for dementia assessment	ZPD-S	ZPD-S	ZPD-S	ZPD-S
Action	**Diagnosis and Intervention**				
Level 0	Finds the information, understands the purpose of the diagnosis and intervention support	A	A	A	A
Level 1	Able to utilise the base for diagnosis-functionality to reason about own assessment	A	A	A	A
Level 2	Familiar with the most common cognitive diseases diagnoses at different levels	ZPD-S	ZPD-S	ZPD-S	A
Level 3	Full understanding of the function of the diagnostic support and its role in reasoning, able to evaluate the outcome of the system in relation to the patient case and own knowledge	ZPD-H	ZPD-H	ZPD-S	A
Action	**Assessment Procedure**				
Level 0	Uses the system as checklist in a linear way and comes to a partial solution, needing to move back and forth to explore essential features	A	A	A	A
Level 1	Uses the system as checklist in a linear way and reaches a final solution	ZPD-S	ZPD-S	A	A
Level 2	Uses the What to do next-functionality and reaches a final solution	ZPD-S	ZPD-S	ZPD-S	A
Level 3	Able to effectively utilize the system as guide, moving back and forth between data capture and diagnosis tabs in an efficient manner	ZPD-S	ZPD-S	ZPD-S	ZPD-S

Figure 9: Summary of the use and learning activity of the users. Person A and B are novices, C has a little experience and D is most experienced physician. "A" means ability to conduct autonomously, "ZPD" means that the action or activity lies in their ZPD, where ZPD-S and ZPD-H mean that the user can be guided by the system or another more experienced human as the more capable peer, respectively.

The analysis method was based on the AAIMA protocol [71] and the concept of a Zone of Proximal Development (ZPD) [111]. An overview of the analysis is shown in Figure 9. From the result, the learning behavior can be identified.

To summarize, qualitative and formative user studies and case studies have been conducted, partly to test different methodologies for collecting or analyzing data. The results are merely indicative since the numbers of users are low and should be further explored in broader user studies. However, the results show that patterns can be detected, and a clear learning behavior is identified in the participating novice user, where the novice is seen to develop towards an expert in the used reasoning pattern after a number of patient cases. As such, the results guide further development of personalization of the educational function of CDSSs.

Chapter 5

Contributions

The contributions specific to the included articles are listed below:

(1) Theory and methods for representing knowledge in the CDSS;

(2) Algorithms and technology for reasoning and decision making (human-agent collaborative reasoning) in the CDSS;

(3) The ACKTUS architecture and system modules for transforming domain knowledge into computer-interpretable guidelines and for developing CDSSs;

(4) ACKTUS modules for designing interaction for supporting humans reasoning and skill development in a CDSS;

(5) Methods for detecting human reasoning, skills, and development of skills in a user when diagnosing dementia using the CDSS DMSS-W.

More details are provided in the following sections.

5.1 Paper I: Hypothesis-Driven Agent Dialogues for Dementia Assessment

***Chunli Yan**, Helena Lindgren. "Hypothesis-Driven Agent Dialogues for Dementia Assessment." In Proceedings of VIII Workshop on Agents Applied in Health Care (A2HC), Murcia, Spain, pp. 13-24, 2013.*

Summary

A wish for a rapid assessment tool was expressed in user studies of earlier versions of DMSS, as a complement to the thorough assessment guided by DMSS and its underlying diagnostic criteria. This request was primarily expressed by clinicians who were not familiar with the dementia diagnosis, and as a way to rapidly test potential hypothetical diagnoses. For this purpose, a hypothesis-driven inquiry dialogue system was developed and implemented in DMSS-W, visualizing the dialogue line and resulting arguments. A goal was to design the dialogue system to allow the clinician to begin with a hypothesis, yet, to guide the clinician to assess enough clinical data about a patient so that a well-founded diagnosis can be suggested. In this process, it was expected that the clinician would learn about dementia assessment, even if jumping to a hypothetical diagnosis.

Another goal was to design the dialogue system following how humans reason, as described in Chapter 4, where novice clinicians tend to use causal reasoning beginning in a hypothesis, and following how humans shift focus between topics in activity, as described by activity theory.

Argument-based dialogues were explored, and the type of inquiry dialogue was found suitable for the purpose. Research on inquiry systems was further developed and adapted to the dementia assessment case. This development included modifying inquiry dialogue systems using defeasible logic with integrated preference levels, and with nested sub-dialogues to follow the natural focus shifts of the human. By applying sub-dialogues, argument evaluation also needed to be conducted and integrated with the dialogue and not as a final step, which is the typical approach in theoretical argumentation dialogue research.

The dialogue system was designed as a MAS with two agents representing the clinician and the domain expert. The clinician selects a tentative hypothesis, and the MAS generates the dialogue with its moves, and the final potential arguments may support or contradict the initial hypothesis. Then it is the clinician who selects the argument that seems most well-founded based on, e.g., the preferred clinical guidelines.

The proposed modified solution was evaluated by comparing its performance to the inquiry dialogue systems presented by Black and Hunter [16]. The proposed solution was shown to be more efficient and applicable in a real use situation. Actually, it represents one of the very few examples of implementation of inquiry dialogues in a real use case.

The results of an initial evaluation study involving domain expert physicians indicate that the approach may be useful to support clinicians in their decision making. However, since the combined MAS and inquiry dialogue system approach is novel and highly unconventional compared to traditional CDSSs, further user studies need to be conducted to evaluate the interaction aspects.

Division of work

HL provided the use case and previous research, while CY suggested the approach and developed and implemented the multi-agent inquiry dialogue system based on defeasible logic and preference levels on sources. The analyses and comparisons with earlier inquiry dialogue systems were conducted by CY. HL conducted the evaluation involving clinicians. CY authored the paper and HL participated in the discussion and revision of the paper.

5.2 Paper II: A Dialogue-Based Approach for Dealing with Uncertain and Conflicting Information in Medical Diagnosis

Chunli Yan, *Helena Lindgren, Juan Carlos Nieves. "A Dialogue-Based Approach for Dealing with Uncertain and Conflicting Information in the Settings of Practical Medical Diagnosis." Accepted by the journal Autonomous Agents and Multi-Agent Systems.*

Summary

Paper II presents the theoretical foundation of an improved version of the inquiry dialogue system presented in Paper I, and the theoretical foundation of the reasoning engine of the DMSS-W, which is included as a module in the extended ACKTUS architecture presented in Paper IV. In the developed theoretical foundation, possibilistic logic was combined with formal argumentation theory where possibilistic logic was used for representing uncertain information and argumentation was used for reasoning with inconsistent and conflicting knowledge. Two algorithms were developed for evaluating the conflicting arguments. The choice of the algorithms depends on the application or the user's preference and need.

The main contributions are (1) a MAS based on possibilistic logic, formal argumentation theory, and a developed version of Black and Hunter's inquiry dialogue systems; (2) the combination of these theories allowed for the transparent generation of potentially conflicting arguments in favor of or against the initially suggested hypothesis; (3) a formal foundation for the representation of knowledge and reasoning with uncertain, inconsistent, and incomplete information; (4) an approach to transparent human-agent collaborative reasoning where the human is allowed to begin with a hypothesis, yet is prevented from missing critical information in the reasoning process; and (5) algorithms applying possibilistic logic that behave in a way that is compliant with the domain knowledge and with what medical experts would expect from a CDSS in the dementia domain.

Division of work

CY proposed and developed the theoretical foundation of an improved version of the inquiry dialogue system combing possibilistic logic with argumentation. CY also developed the algorithms and authored the paper. HL and JCN participated in the discussion about how to present the MAS, and HL evaluated the algorithms from a medical perspective using medical examples. HL and JCN also suggested revisions of the paper.

5.3 Paper III: ACKTUS - A Platform for Developing Personalized Support Systems in the Health Domain

*Helena Lindgren, **Chunli Yan**. "ACKTUS - A Platform for Developing Personalized Support Systems in the Health Domain." In Proceedings of the 5th International Conference on Digital Health 2015, pp. 135-142, 2015.*

Summary

Paper III presents ACKTUS, a Semantic Web platform for modeling and managing knowledge integrated into medical applications and for designing the interaction with the end user. A core ontology, implemented using OWL/RDF and stored in Sesame, serves as the information model and KB in the system. It is based on the ICF developed by the World Health Organization and integrates a version of the AIF modified to capture uncertain information. The ontology also includes a part for building the structure of each user interface as well as for defining the overall activity to be supported by the resulting CDSS. Additional ontologies were also developed for representing information about the actor and events occurring when actors are using an ACKTUS application. A CMS was developed for allowing domain experts to manage and extend the core ontology for knowledge engineering and interaction design purposes for their different projects.

ACKTUS was used for developing DMSS-W, where domain experts were involved in modeling the content of the system presented in earlier studies (e.g., [71]). ACKTUS also serves as the research infrastructure for the development of knowledge-based interventions in the domain of fall prevention in older adults [70], and chronic obstructive pulmonary disorder [82]. User studies have shown that ACKTUS can be used by domain professionals not familiar with knowledge engineering tasks.

Division of work

HL proposed the approach of applying a core ontology and a CMS module for modeling and managing knowledge by the domain expert and designed and developed the core ontology in collaboration with CY. CY developed the CMS module, designed and developed the additional ontologies and implemented the modules. HL authored the main parts of the paper, and CY contributed with a presentation of the modules and participated in revisions of the paper.

5.4 Paper IV: A Generic Approach for Data Management and End-User Development of Clinical Decision Support Systems

Paper IV: Chunli Yan, *Helena Lindgren. "A Generic Approach for Data Management and End-User Development of Clinical Decision Support Systems." Technical report / UMINF 18.08, ISSN 0348-0542, Umeå University, Umeå, 2018.*

Summary

Paper IV presents the design of the support for the different reasoning strategies introduced in Chapter 4, and the technology developed to implement the interactive reasoning functionality. The design and structure of the user interface play an important role in providing reasoning support and, therefore, an ontology-driven method was selected. Further, to facilitate knowledge elicitation and end-user development, the ACKTUS architecture presented in Paper III was extended with two additional modules: (1) an ontology-driven GUI generator that automatically generated the user interface whenever the user logs in; (2) an inference engine supporting two reasoning strategies. These modules were used in the development of DMSS-W and can be reused when developing new ACKTUS-based CDSSs.

The GUI generator generates the user interface based on the activity theory model of purposeful activity embedded in the ontology. Three particular assessment protocols are key in the GUI generation, one that defines the activity aimed to be supported by the application, one that defines the data capture sub-activity, and one that defines the reasoning and decision making related to diagnosis and intervention sub-activity.

The inference engine handles uncertainty and generates diagnoses by implementing the possibilistic framework presented in Paper II. The engine extracts the knowledge modeled by domain experts using the structures based on AIF, combined with patient information collected using the CDSS, into possibilistic information.

The different ways that medical professionals conduct clinical reasoning and decision making presented in Chapter 4 formed the foundation of the design and development of the DMSS-W. It was developed for supporting clinical reasoning and decision making and continuing medical education in the users of the system. A pilot study of DMSS-W was conducted and presented in the paper, involving four medical professionals with different levels of expertise. Results indicate that the strategies are complementary and serve different purposes, and can support users with varying levels of experience and skills.

Division of work

CY proposed the methods for dynamically generating a CDSS and designed, developed, and integrated the new modules in ACKTUS. Further, CY implemented the CDSS for dementia, based on ACKTUS, and monitored the evaluation studies. CY authored the paper and HL participated in the discussions and took part in revising the paper.

5.5 Paper V: Detecting Learning and Reasoning Patterns in a CDSS for Dementia Investigation

Paper V: *Helena Lindgren,* ***Chunli Yan****. "Detecting Learning and Reasoning Patterns in a CDSS for Dementia Investigation." Studies in Health Technology and Informatics 210: 739-742, 2015.*

Summary

An earlier study of DMSS showed differences in how clinicians complied with the suggestions provided by the CDSS [67]. A qualitative analysis was done, which indicated that a lack of knowledge could be the reason for non-compliance. It was further investigated whether individual clinicians' reasoning and skill development could be detected through reasoning patterns discovered when they use DMSS-W in clinical practice. If this information could be tracked, then personally-tailored support for continuing medical education could be developed.

Reasoning conducted in clinical practice is manifested through different and often combined reasoning and learning strategies, as described in Chapter 4. The purpose of the study presented in Paper V was to develop a method for detecting reasoning patterns and skill levels. This purpose was achieved by taking the activity-theoretical design of the ontology-generated user interface as the starting point and developing an event logging functionality based on the ontology.

A CDSS application such as DMSS-W is designed using activity theory models of purposeful activity (Chapter 3). The main assessment protocol contains an ordered list of nested sub-protocols, which builds the menu system in the user interface. In practice, the protocols represent a hierarchy of sub-actions in the assessment task with particular goals (see Figure 8). As a consequence, each click event is interpreted as manifesting a particular goal and having a particular purpose.

An information model for the log data was developed, where each event was logged with information about the time, actor, type of activity, and if evidence was generated. Four users were selected for a case study that represented two novices and two experts on dementia. Their clicking behaviors were used to analyze their reasoning and learning strategies. The results showed that a distinction could be made between the two types of users included in this case study.

Division of work

HL proposed the approach and CY designed the method to detect learning and reasoning patterns, based on earlier studies and theory, and developed the information model, the repository for the log data, and implemented the log functionality. HL authored the paper in dialogue with CY, and CY contributed with revisions.

5.6 Paper VI: Diagnostic Reasoning Guided by a Decision-Support System: a Case Study

Chunli Yan, Helena Lindgren. "Diagnostic Reasoning Guided by a Decision-Support System: a Case Study." In Proceedings of the ACM European Conference on Cognitive Ergonomics (ECCE-17), Umeå, pp. 25-30, 2017.

Summary

The purpose of the study presented in Paper VI was to develop a method for detecting the potential development of skills and knowledge in a user of DMSS-W. This research was conducted by studying log data from users over a longer period of time and a number of completed patient cases. The key research question addressed was whether learning and skill development can be detected.

A semi-automatic analysis of the log file was developed and implemented that categorized logs into three types of sessions: (1) the patient case is complete; (2) partial information is entered about a patient case; and (3) no patient-specific information is entered, meaning the user navigates around the user interface without assessing a patient.

A case study was conducted where a physician, who was a novice to both the application and the diagnosis of dementia, was studied and compared to the case of an expert physician using the system. Differences between them were found, and a clear pattern indicating that learning took place, both regarding how to use the system and how to diagnose dementia, was observed in the novice user. Reasoning patterns were detected and analyzed based on the theoretical base presented in Chapter 4.

The result serves as a starting point for further study, where more users would be involved.

Division of work

CY proposed the method to detect reasoning patterns, analyzed log data, conducted the case study, and authored the paper in dialogue with HL.

Bibliography

[1] Leila Amgoud and Claudette Cayrol. A reasoning model based on the production of acceptable arguments. *Annals of Mathematics and Artificial Intelligence*, 34:197–215, 2002.

[2] Leila Amgoud, Claudette Cayrol, Marie-Christine Lagasquie-Schiex, and Pierre Livet. On bipolarity in argumentation frameworks. *International Journal of Intelligent Systems*, 23(10):1062–1093, 2008.

[3] Leila Amgoud, Nicolas Maudet, and Simon Parsons. Modelling dialogues using argumentation. In *Proceedings of the 4th Conference on Multi-Agent Systems*, pages 31–38. IEEE, 2000.

[4] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider. *The description logic handbook: Theory, implementation and applications.* Cambridge university press, 2003.

[5] Franz Baader and Ulrike Sattler. An overview of tableau algorithms for description logics. *Studia Logica*, 69(1):5–40, 2001.

[6] Jakob E. Bardram. A novel approach for creating activity-aware applications in a hospital environment. In *IFIP Conference on Human-Computer Interaction*, pages 731–744. Springer, 2009.

[7] Jakob E. Bardram and Thomas R. Hansen. Why the plan doesn't hold - a study of situated planning, articulation and coordination work in a surgical ward. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*, pages 331–340. ACM, 2010.

[8] David Bassen, Saurabh Nayak, Xia Chong Li, Mitchell Sam, Jagmohan Sidhu, Martha F Nelson, and Walker H Land. Clinical decision support system (cdss) for the classification of atypical cells in pleural effusions. *Procedia Computer Science*, 20:379–384, 2013.

[9] Montserrat Batet, Karina Gibert, and Aida Valls. The data abstraction layer as knowledge provider for a medical multi-agent system. In *Knowledge Management for Health Care Procedures*, pages 87–100. Springer, 2008.

[10] Patrick Emanuel Beeler, David Westfall Bates, and Balthasar Luzius Hug. Clinical decision support systems. *Swiss Medical Weekly*, 144:w14073, 2014.

[11] J. H. van Bemmel and M. A. Musen. *Handbook of Medical Informatics.* Heidelberg: Springer-Verlag, 1997.

[12] Trevor J. M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.

[13] Trevor J. M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171:619–641, 2007.

[14] Casey C. Bennett and Kris Hauser. Artificial intelligence framework for simulating clinical decision-making: A markov decision process approach. *Artificial Intelligence in Medicine*, 57(1):9–19, 2013.

[15] Tim Berners-Lee, James Hendler, Ora Lassila, et al. The semantic web. *Scientific American*, 284(5):34–43, 2001.

[16] Elisabeth Black and Antony Hunter. An inquiry dialogue system. *Autonomous Agents and Multi-Agent Systems*, 19(2):173–209, 2009.

[17] Elizabeth Black and Anthony Hunter. A generative inquiry dialogue system. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, page 241. ACM, 2007.

[18] Alessio Bottrighi and Paolo Terenziani. Meta-glare: A meta-system for defining your own computer interpretable guideline system - architecture and acquisition. *Artificial Intelligence in Medicine*, 72:22–41, 2016.

[19] Andrew E. Budson and Paul R. Solomon. *Memory Loss: A Practical Guide for Clinicians.* Elsevier, 2011.

[20] A. Burns and S. Iliffe. Dementia. *BMJ*, 338:b75, 2009.

[21] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Daniele Nardi. Reasoning in expressive description logics. In Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, pages 1581–1634. Elsevier Science Publishers, 2001.

[22] Sougata Chakraborty and Shibakali Gupta. Medical application using multi agent system-a literature survey. *International Journal of Engineering Research and Applications*, 4(2):528–546, 2014.

[23] Carlos Chesñevar, Jarred , Mcginnis, Sanjay Modgil, Iyad Rahwan, Chris Reed, Guillermo Simari, Matthew South, Gerard Vreeswijk, Steven Willmott, et al. Towards an argument interchange format. *The Knowledge Engineering Review*, 21(4):293–316, 2006.

[24] Carlos I. Chesñevar, Guillermo R. Simari, Teresa Alsinet, and Lluís Godo. A logic programming framework for possibilistic argumentation with vague knowledge. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pages 76–84. AUAI Press, 2004.

[25] Betty Collis and Anoush Margaryan. Applying activity theory to computer-supported collaborative learning and work-based activities in corporate settings. *Educational Technology Research and Development*, 52(4):38–52, 2004.

[26] Ulises Cortés, Roberta Annicchiarico, Javier Vázquez-Salceda, Cristina Urdiales, Lola Cañamero, Maite López, Miquel Sànchez-Marrè, and Carlo Caltagirone. Assistive technologies for the disabled and for the new generation of senior citizens: the e-tools architecture. *AI Communications*, 16(3):193–207, 2003.

[27] Jiangbo Dang, Amir Hedayati, Ken Hampel, and Candemir Toklu. An ontological knowledge framework for adaptive medical workflow. *Journal of Biomedical Informatics*, 41(5):829–836, 2008.

[28] Gil Dekel. *Inspiration: a functional approach to creative practice.* PhD thesis, University of Portsmouth, 2009.

[29] D. Dubois, J. Lang, and H. Prade. Possibilistic logic. In *Handbook of logic in artificial intelligence and logic programming (vol. 3)*, pages 439–513. Oxford University Press, Inc., 1994.

[30] Didier Dubois and Henri Prade. Possibilistic logic: a retrospective and prospective view. *Fuzzy Sets and Systems*, 144(1):3–23, 2004.

[31] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.

[32] Phan Minh Dung, Robert A. Kowalski, and Francesca Toni. Dialectic proof procedures for assumption-based, admissible argumentation. *Artificial Intelligence*, 170(2):114–159, 2006.

[33] Claudio Eccher, Antonella Ferro, Andreas Seyfang, Marco Rospocher, and Silvia Miksch. Modeling clinical protocols using semantic mediawiki: The case of the oncocure project. In *ECAI Workshop on Knowledge Management for Health Care Procedures (K4HelP)*, pages 42–54. Springer, 2008.

[34] June S. Eichner, Maya Das, et al. *Challenges and barriers to clinical decision support (CDS) design and implementation experienced in the Agency for Healthcare Research and Quality CDS demonstrations.* Agency for Healthcare Research and Quality Rockville, 2010.

[35] David A. Evans and Vilma L. Patel. *Cognitive science in medicine: biomedical modeling.* MIT Press, 1989.

[36] John Fox, Nicky Johns, and Ali Rahmanzadeh. Disseminating medical knowledge: the PROforma approach. *Artificial Intelligence in Medicine*, 14(1):157–181, 1998.

[37] John Fox, Nicky Johns, Ali Rahmanzadeh, and Richard Thomson. PROforma: A method and language for specifying clinical guidelines and protocols. In *Proceedings of Medical Informatics Europe*, pages 516–520. IOS Press, 1996.

[38] Elizabeth Furtado, João José Vasco Furtado, Wilker Bezerra Silva, Daniel William Tavares Rodrigues, Leandro da Silva Taddeo, Quentin Limbourg, and Jean Vanderdonckt. An ontology-based method for universal design of user interfaces. In *Task Models and Diagrams for User Interface Design (TAMODIA 2002)*, 2002.

[39] Elena García-Barriocanal, Miguel-Angel Sicilia, and Salvador Sánchez-Alonso. Usability evaluation of ontology editors. *Knowledge Organization*, 32(1):1–9, 2005.

[40] Amit X. Garg, Neill K. J. Adhikari, Heather McDonald, M. Patricia Rosas-Arellano, P. J. Devereaux, Joseph Beyene, Justina Sam, and R. Brian Haynes. Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: a systematic review. *JAMA*, 293(10):1223–1238, 2005.

[41] Sébastien Gebus and Kauko Leiviskä. Knowledge acquisition for decision support systems on an electronic assembly line. *Expert Systems with Applications*, 36(1):93–101, 2009.

[42] John H. Gennari, Mark A. Musen, Ray W. Fergerson, William E. Grosso, Monica Crubézy, Henrik Eriksson, Natalya F. Noy, and Samson W. Tu. The evolution of Protégé: an environment for knowledge-based systems development. *International Journal of Human-Computer Studies*, 58(1):89–123, 2003.

[43] Ignasi Gómez-Sebastià, Jonathan Moreno, Sergio Álvarez-Napagao, Dario Garcia-Gasulla, Cristian Barrué, and Ulises Cortés. Situated agents and humans in social interaction for elderly healthcare: From coaalas to AVICENA. *Journal of Medical Systems*, 40(2):38, 2016.

[44] Horacio González-Vélez, Mariola Mier, Margarida Julià-Sapé, Theodoros N Arvanitis, Juan M García-Gómez, Montserrat Robles, Paul H Lewis, Srinandan Dasmahapatra, David Dupplaw, Andrew Peet, et al. Healthagents: distributed multi-agent brain tumor diagnosis and prognosis. *Applied Intelligence*, 30(3):191–202, 2009.

[45] N. L. Griffin and F. D. Lewis. A rule-based inference engine which is optimal and vlsi implementable. In *IEEE International Workshop on Tools for Artificial Intelligence. Architectures, Languages and Algorithms*, pages 246–251. IEEE, 1989.

[46] Thomas R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.

[47] Hans-Jörg Happel, Axel Korthaus, Stefan Seedorf, and Peter Tomczyk. Kontor: an ontology-enabled approach to software reuse. In *Proceedings of the Eighteenth International Conference on Software Engineering and Knowledge Engineering (SEKE)*, 2006.

[48] Ahnaf Rashik Hassan and Abdulhamit Subasi. A decision support system for automated identification of sleep stages from single-channel EEG signals. *Knowledge-Based Systems*, 128:115–124, 2017.

[49] Barbara Hayes-Roth. An architecture for adaptive intelligent systems. *Artificial Intelligence*, 72(1-2):329–365, 1995.

[50] Benjamin Heitmann, Sheila Kinsella, Conor Hayes, and Stefan Decker. Implementing semantic web applications: reference architecture and challenges. *Proceedings of the 5th Workshop on Semantic Web Enabled Software Engineering.*, 524, 2009.

[51] Earl Hunt. Cognitive science: Definition, status, and questions. *Annual Review of Psychology*, 40(1):603–629, 1989.

[52] David Isern, David Sánchez, and Antonio Moreno. Agents applied in health care: A review. *International Journal of Medical Informatics*, 79(3):145–166, 2010.

[53] Peter D. Johnson, Samson Tu, Nick Booth, Bob Sugden, and Ian N. Purves. Using scenarios in chronic disease management guidelines for primary care. In *Proceedings of the AMIA Symposium*, pages 389–393. American Medical Informatics Association, 2000.

[54] Björn A. Johnsson and Boris Magnusson. Towards end-user development of graphical user interfaces for internet of things. *Future Generation Computer Systems*, in press.

[55] Victor Kaptelinin. Learning with artefacts: integrating technologies into activities. *Interacting with Computers*, 15(6):831–836, 2003.

[56] Victor Kaptelinin and Michael Cole. Individual and collective activities in educational computer game playing. In *Proceedings of the 2nd International Conference on Computer Support for Collaborative Learning*, pages 142–147. International Society of the Learning Sciences, 1997.

[57] Daniel Karlsson. *Aspects of the use of medical decision-support systems: the role of context in decision support.* PhD thesis, Linköping University, 2001.

[58] Nishan C. Karunatillake, Nicholas R Jennings, Iyad Rahwan, and Sarvapali D Ramchurn. Managing social influences through argumentation-based negotiation. In *Argumentation in Multi-Agent Systems*, pages 107–127. Springer, 2007.

[59] Alison L. Kidd. *Knowledge acquisition for expert systems: A practical handbook.* Plenum Press, New York, 1987.

[60] Dongmin Kim and Izak Benbasat. Trust-related arguments in internet stores: A framework for evaluation. *Journal of Electronic Commerce Research*, 4(2):49–64, 2003.

[61] Vassilis G. Koutkias, Ioanna Chouvarda, and Nicos Maglaveras. A multiagent system enhancing home-care health services for chronic disease management. *IEEE Transactions on Information Technology in Biomedicine*, 9(4):528–537, 2005.

[62] Jérôome Lang. Possibilistic logic as a logical framework for min-max discrete optimisation problems and prioritized constraints. In *Fundamentals of Artificial Intelligence Research*, pages 112–126. Springer, 1991.

[63] Jill Larkin, John McDermott, Dorothea P. Simon, and Herbert A. Simon. Expert and novice performance in solving physics problems. *Science*, 208(4450):1335–1342, 1980.

[64] A. Latoszek-Berendsen, H. Tange, H. J. Van Den Herik, A. Hasman, et al. From clinical practice guidelines to computer-interpretable guidelines. *Methods of Information in Medicine*, 49(6):550–570, 2010.

[65] Yuangui Lei, Enrico Motta, and John Domingue. Design of customized web applications with OntoWeaver. In *Proceedings of the 2nd International Conference on Knowledge Capture*, pages 54–61. ACM, 2003.

[66] Helena Lindgren. Integrating clinical decision support system development into a development process of clinical practice – experiences from dementia care. In M. Peleg, N. Lavrac, and C. Combi, editors, *Artificial Intelligence in Medicine (AIME 2011)*, volume 6747 of *Lecture Notes in Computer Science*, pages 129–138. Springer, 2011.

[67] Helena Lindgren. Limitations in physicians knowledge when assessing dementia diseases - an evaluation study of a decision-support system. *Studies in Health Technology and Informatics*, 169:120–124, 2011.

[68] Helena Lindgren, Lage Burström, and Bengt Järvholm. Developing ambient support technology for risk management in the mining industry. In Carlos Ramos, Paulo Novais, Celine Ehrwein Nihan, and Juan M. Corchado Rodriguez, editors, *Ambient Intelligence – Software and Applications*, volume 291 of *Advances in Intelligent Systems and Computing*, pages 161–169. Springer International Publishing, 2014.

[69] Helena Lindgren, Ming-Hsin Lu, Yeji Hong, and Chunli Yan. Applying the zone of proximal development when evaluating clinical decision support systems: a case study. *Studies in Health Technology and Informatics*, 247:131–135, 2018.

[70] Helena Lindgren, Lillemor Lundin-Olsson, Petra Pohl, and Marlene Sandlund. End users transforming experiences into formal information and process models for personalised health interventions. *Studies in Health Technology and Informatics*, 205:378–382, 2014.

[71] Helena Lindgren, Patrik J. Winnberg, and Peter Winnberg. Domain experts tailoring interaction to users - an evaluation study. In Pedro Campos, T. C. Nicholas Graham, Joaquim A. Jorge, Nuno Jardim Nunes, Philippe A. Palanque, and Marco Winckler, editors, *INTERACT (3)*, volume 6948 of *Lecture Notes in Computer Science*, pages 644–661. Springer, 2011.

[72] Helena Lindgren, Patrik J. Winnberg, and Chunli Yan. Collaborative development of knowledge-based support systems: A case study. *Studies in Health Technology and Informatics*, 180:1111–1113, 2012.

[73] Pattie Maes. Artificial life meets entertainment: lifelike autonomous agents. *Communications of the ACM*, 38(11):108–114, 1995.

[74] Catherine C. Marinagi, Constantine D. Spyropoulos, Christos Papatheodorou, and Stavros Kokkotos. Continual planning and scheduling for managing patient tests in hospital laboratories. *Artificial Intelligence in Medicine*, 20(2):139–154, 2000.

[75] Peter McBurney, Rogier M. Van Eijk, Simon Parsons, and Leila Amgoud. A dialogue-game protocol for agent purchase negotiations. *Autonomous Agents and Multi-Agent Systems*, 7(3):235–273, 2003.

[76] Peter McBurney and Simon Parsons. Representing epistemic uncertainty by means of dialectical argumentation. *Annals of Mathematics and Artificial Intelligence*, 32(1-4):125–169, 2001.

[77] Peter McBurney and Simon Parsons. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information*, 11(3):315–334, 2002.

[78] Silvia Miksch, Yuval Shahar, and Peter Johnson. Asbru: a task-specific, intention-based, and time-oriented language for representing skeletal plans. In *Proceedings of the 7th Workshop on Knowledge Engineering: Methods and Languages (KEML-97)*, pages 9.1–9.20. The Open University, Milton Keynes, UK, 1997.

[79] Annemarie Mol. *The body multiple: Ontology in medical practice*. Duke University Press, 2002.

[80] Mark A. Musen, Blackford Middleton, and Robert A. Greenes. Clinical decision-support systems. In E. Shortliffe and J. Cimino, editors, *Biomedical Informatics*, pages 643–674. Springer, London, 2014.

[81] Brad A. Myers and Mary Beth Rosson. Survey on user interface programming. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 195–202. ACM, 1992.

[82] André Nyberg, Karin Wadell, Helena Lindgren, and Malin Tistad. Internet-based support for self-management strategies for people with COPD-protocol for a controlled pragmatic pilot trial of effectiveness and a process evaluation in primary healthcare. *BMJ Open*, 7(7), 2017.

[83] Jerome A. Osheroff, Jonathan M. Teich, Blackford Middleton, Elaine B. Steen, Adam Wright, and Don E. Detmer. A roadmap for national action on clinical decision support. *Journal of the American Medical Informatics Association: JAMIA*, 14(2):141–145, 2007.

[84] Simon Parsons, Michael Wooldridge, and Leila Amgoud. An analysis of formal inter-agent dialogues. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 1*, pages 394–401. ACM, 2002.

[85] Vimla L. Patel, Vanessa G. Allen, José F. Arocha, and Edward H. Shortliffe. Representing clinical guidelines in GLIF: individual and collaborative expertise. *Journal of the American Medical Informatics Association*, 5(5):467–483, 1998.

[86] Vimla L. Patel and Guy J. Groen. Knowledge based solution strategies in medical reasoning. *Cognitive Science*, 10(1):91–116, 1986.

[87] Vimla L. Patel and Guy J. Groen. Developmental accounts of the transition from medical student to doctor: some problems and suggestions. *Medical Education*, 25(6):527–535, 1991.

[88] Vimla L. Patel and Guy J. Groen. The general and specific nature of medical expertise: A critical look. In A. Ericsson and J. Smith, editors, *Toward a general theory of expertise: Prospects and limits*, pages 93–125. Cambridge University Press, 1991.

[89] Vimla L. Patel, David R. Kaufman, and Jose F. Arocha. Emerging paradigms of cognition in medical decision-making. *Journal of Biomedical Informatics*, 35(1):52–75, 2002.

[90] Vimla L. Patel, David R. Kaufman, and Sheldon A. Magder. The acquisition of medical expertise in complex dynamic environments. In A. Ericsson, editor, *The Road to Excellence: The Acquisition of Expert Performance in the Arts and Sciences, Sports and Games*, pages 127–165. Lawrence Erlbaum Mahwah, NJ, 1996.

[91] Heiko Paulheim and Florian Probst. Ontology-enhanced user interfaces: A survey. *International Journal on Semantic Web and Information Systems (IJSWIS)*, 6(2):35–59, 2010.

[92] Mor Peleg. Computer-interpretable clinical guidelines: a methodological review. *Journal of Biomedical Informatics*, 46(4):744–763, 2013.

[93] Mor Peleg, Aziz A. Boxwala, Omolola Ogunyemi, Qing Zeng, Samson Tu, Ronilda Lacson, Elmer Bernstam, Nachman Ash, Peter Mork, Lucila Ohno-Machado, et al. Glif3: the evolution of a guideline representation format. In *Proceedings of the AMIA Symposium*, pages 645–649. American Medical Informatics Association, 2000.

[94] Mor Peleg, Samson Tu, Jonathan Bury, Paolo Ciccarese, John Fox, Robert A. Greenes, Richard Hall, Peter D. Johnson, Neill Jones, Anand Kumar, et al. Comparing computer-interpretable guideline models: a case-study approach. *Journal of the American Medical Informatics Association*, 10(1):52–68, 2003.

[95] Henry Prakken. Formal systems for persuasion dialogue. *The Knowledge Engineering Review*, 21(02):163–188, 2006.

[96] Silvana Quaglini, Mario Stefanelli, Giordano Lanzola, Vincenzo Caporusso, and Silvia Panzarasa. Flexible guideline-based patient careflow systems. *Artificial Intelligence in Medicine*, 22(1):65–80, 2001.

[97] David Riaño, Francis Real, Joan Albert López-Vallverdú, Fabio Campana, Sara Ercolani, Patrizia Mecocci, Roberta Annicchiarico, and Carlo Caltagirone. An ontology-based personalization of health-care knowledge to support clinical decisions for chronically ill patients. *Journal of Biomedical Informatics*, 45(3):429–446, 2012.

[98] David R. Russell. Looking beyond the interface: Activity theory and distributed learning. *Distributed Learning: Social and Cultural Approaches to Practice*, pages 64–82, 2002.

[99] Stuart Jonathan Russell and Peter Norvig. *Artificial intelligence: a modern approach (Second Edition)*. Prentice hall, 2003.

[100] Andreas Seyfang, Silvia Miksch, Mar Marcos, Jolanda Wittenberg, Cristina Polo-Conde, and Kitty Rosenbrand. Bridging the gap between informal and formal guideline representations. In *ECAI 2006: 17th European Conference on Artificial Intelligence*, volume 141, pages 447–451. IOS Press, 2006.

[101] Yuval Shahar, Silvia Miksch, and Peter Johnson. The asgaard project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines. *Artificial Intelligence in Medicine*, 14(1-2):29–51, 1998.

[102] Erez Shalom, Yuval Shahar, Meirav Taieb-Maimon, Guy Bar, Susana B Martins, Ohad Young, Laszlo Vaszar, Yair Liel, Avi Yarkoni, Mary K Goldstein, et al. Can physicians structure clinical guidelines? experiments with a mark-up-process methodology. In *ECAI workshop on Knowledge Management for Health Care Procedures (K4HelP)*, pages 67–80. Springer, 2008.

[103] Ida Sim, Paul Gorman, Robert A. Greenes, R. Brian Haynes, Bonnie Kaplan, Harold Lehmann, and Paul C. Tang. Clinical decision support systems for the practice of evidence-based medicine. *Journal of the American Medical Informatics Association*, 8(6):527–534, 2001.

[104] Shailendra Singh, Bukhary Ikhwan Ismail, Fazilah Haron, and Chan Huah Yong. Architecture of agent-based healthcare intelligent assistant on grid environment. In *Parallel and Distributed Computing: Applications and Technologies*, pages 58–61. Springer, 2005.

[105] Kênia Sousa. Model-driven approach for user interface: business alignment. In *Proceedings of the 1st ACM SIGCHI symposium on Engineering Interactive Computing Systems*, pages 325–328. ACM, 2009.

[106] Peter Stone and Manuela Veloso. Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3):345–383, 2000.

[107] Alberto Tablado, Arantza Illarramendi, Miren I Bagüés, Jesús Bermúdez, and Alfredo Goni. Aingeru: an innovating system for tele assistance of elderly people. In *TELECARE*, pages 27–36, 2004.

[108] Monica Tentori, Jesus Favela, and Marcela D Rodriguez. Privacy-aware autonomous agents for pervasive healthcare. *Intelligent Systems, IEEE*, 21(6):55–62, 2006.

[109] Samson W. Tu and Mark A. Musen. A flexible approach to guideline modeling. In *Proceedings of the AMIA symposium*, pages 420–424. American Medical Informatics Association, 1999.

[110] Aida Valls, Karina Gibert, David Sánchez, and Montserrat Batet. Using ontologies for structuring organizational knowledge in home care assistance. *International Journal of Medical Informatics*, 79(5):370–387, 2010.

[111] Lev Semenovich Vygotsky. *Mind in society: The development of higher psychological processes.* Cambridge: Harvard university press, 1978.

[112] Douglas N. Walton. *The new dialectic: Conversational contexts of argument.* University of Toronto Press, 1998.

[113] Michael Wooldridge. *An introduction to MultiAgent Systems.* John Wiley & Sons, 2009.

[114] Michael Wooldridge and Nicholas R. Jennings. Intelligent agents: theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.

[115] Chunli Yan, Juan Carlos Nieves, and Helena Lindgren. A multi-agent system for nested inquiry dialogues. In Yves Demazeau, Franco Zambonelli, Juan M. Corchado, and Javier Bajo, editors, *Advances in Practical Applications of Heterogeneous Multi-Agent Systems. The PAAMS Collection. PAAMS 2014. Lecture Notes in Computer Science*, volume 8473, pages 303–314. Springer, Cham, 2014.