

The Doctoral Consortium

At the
Twenty-Third International Conference on
Case-Based Reasoning
(ICCBR 2015)

Frankfurt, Germany
September 2015

Nirmalie Wiratunga and Sarah Jane Delany (Editors)

Program Chairs

Nirmalie Wiratunga	The Robert Gordon University, Scotland
Sarah Jane Delany	Dublin Institute of Technology, Ireland

Program Committee

David W. Aha	Naval Research Laboratory, USA
Klaus-Dieter Althoff	DFKI / University of Hildesheim, Germany
Kerstin Bach	NTNU, Norway
Amélie Cordier	LIRIS, France
Susan Crow	The Robert Gordon University, Scotland
Odd Erik Gundersen	NTNU, Norway
Joseph Kendall-Morwick	Capital University, USA
David Leake	Indiana University, USA
Jean Lieber	LORIA - INRIA Lorraine, France
Cindy Marling	Ohio University, USA
Stewart Massie	The Robert Gordon University, Scotland
Stefania Montani	University Piemonte Orientale, Italy
Miltos Petridis	University of Brighton, UK
Barry Smyth	University College Dublin, Ireland
David Wilson	University of North Carolina at Charlotte, USA

Preface

The Doctoral Consortium volume contains the research summaries that were presented at the 7th Annual ICCBR 2015 Doctoral Consortium (www.iccbr15.de/index.php/2014-10-27-10-03-12/doctoral-consortium) held on 28th Sep 2015 in Frankfurt, Germany. There were 11 accepted submissions consisting of: (1) a 1-page Application Cover Page; (2) a 3-page Research Summary; (3) a 1-page Resum; and (4) a letter of support from the student's advisor. The objectives, progress, plans and references in each research summary have been progressively refined according to feedback from one or more PC members, of which one was also the assigned mentor. Feedback was organised into three broad areas: general outlook in terms of research hypothesis and proposed methodology; detailed comments specific to the student's project; and finally advice for the talk presentation. A face-to-face pre-event meeting opportunity enabled all student-mentor pairs to refine their presentations. Mentors also had the responsibility of leading the question and answer session following each mentee presentations on the day.

The ICCBR-15 DC began on September 27th with an informal meet and greet session, followed by a discussion led by Dr Kerstin Bach (Norwegian University of Science and Technology) on shared student experiences. The evening ended with dinner sponsored by the conference. On September 28th, the formal program started with an invited talk by Prof Klaus-Dieter Althoff (Hildesheim University), entitled Lessons Learned - A Journey from Research Student to Professor. The rest of the program consisted of 15-minute talks presented by 11 doctoral students on their Research Summary. The presentations covered a wide range of CBR topics including recommender systems, retrieval, adaptation and maintenance in CBR, e-learning systems, agents and analogical reasoning, distributed CBR and AI in music.

Many people participated in making the DC event a success. We wish to thank all our PC members who provided important and useful guidance to DC students, either as reviewers or as mentors. We are very grateful for the generous support of the sponsors of the ICCBR-15 DC: The AI Journal and National Science Foundation. Once again AIJ has enabled us to provide significantly discounted registration fees to our participants and the NSF funding obtained through David Wilson has helped fund travel costs for our students from the US. Finally thanks go out to David W. Aha who has helped muster a healthy number of participants for this year's event.

Finally thank you to all our DC participants. We trust that the ICCBR-15 DC enhanced your interest in studying CBR and that the welcome and support from the CBR community has sparked your interest in this field for many years to come.

September 2015
Frankfurt

Nirmalie Wiratunga
Sarah Jane Delany

Distributed Case-based Support for the Architectural Conceptualization Phase

Viktor Ayzenshtadt

University of Hildesheim, Institute of Computer Science
 Competence Center Case-Based Reasoning
 German Research Center for Artificial Intelligence
 Trippstadter Straße 122, 67663 Kaiserslautern, Germany
 ayzensht@uni-hildesheim.de

1 Introduction

When an architect conceptualizes a new building she is very likely in need of new ideas, solutions and inspiration to create a new design. *Metis* [3] is a basic case-based design research project of the German Research Center for Artificial Intelligence (DFKI) and the KSD Research Group of the TU Munich that aims to help architects during the early design concept stage and corresponding building plans creation by providing them with similar building designs to a created one. One of the main aspects of the project is the creation of new cases (building designs) by transforming floorplan sketches with image processing techniques into graph representations which are based on the *Semantic Fingerprint* [4] model. Another one is the retrieval process that uses a multi-agent system with case-based agents that are able to apply either subgraph matching or CBR-framework-based retrieval to find similar building designs. An architect can search for them by using a web browser-based graphical interface. As usual, the project also includes participation of experts, who discuss and explain the details and aspects of the CAAD, CBR and Multi-agent systems research tasks.

In my master thesis I extended the previously existing initial concept of the retrieval system to provide the core functionality for the project's retrieval tasks. This system uses the retrieval container structure where each container acts as a separate multi-agent system that is only responsible for resolving a single user query. The retrieval process is coordinated by a corresponding agent. The case base consists of extracted and imported graph representations of the building designs. The gateway supports the connection between the core systems and the user interface.

2 PhD Research Focus

In my PhD thesis research I am going to concentrate on the research fields named in Section 1 and continue to study the case-based architectural design support questions. The implemented retrieval system from the master thesis will be taken as a base and extended for the further research. In detail, the currently planned research goals are described in the following sections.

2.1 Case Representation

This research part will answer the question which model is the most preferable one for representing architectural design cases in CBD applications – graphs or *attribute-value* concepts. The comparison of those models will include the study on how both of them perform under the same conditions when conducting retrieval and inserting of new cases. Currently cases consist of graph-based, GraphML-based [1], myCBR-based [2] and ontologically applied multi-agent communication language FIPA-SL-based floorplan representations that include room representations and room connections with corresponding attributes and values. The knowledge for creating those cases is acquired and maintained by the specific maintainer system agent that obtains, transforms, separates and inserts building design graphs into the corresponding case bases.

2.2 Retrieval Performance

The cross-validation of both retrieval approaches – *subgraph matching* and *CBR-framework-based* – is another part of the planned research. Here both approaches will be validated by applying the cross-comparison between those two types. The aim of this process is to answer the question which of both approaches provides the highest quality of the retrieval results. Both retrieval models will be confronted with different user scenarios to find the best suitable method for a given situation or context.

2.3 Retrieval Coordination

Two currently available retrieval coordination approaches – *rule-based* and *case-based* – are going to be extended to a full functionality and provide a complete pool of features needed for the relevant query. In addition a cross-comparison of them as a part of the retrieval performance measurement could be performed as well. Architectural experts' help and users' feedback can be taken into account and used for the evaluation of the result quality.

2.4 CBR Domain Modelling

The myCBR part of the retrieval system contains the CBR domain that is based on the structure of the *Semantic Fingerprint* model. The underlying model of the domain is going to be improved (with the experts' help inter alia) and adapted to the results of the studies named in the previous research goals.

This aim is also valid for the CBR agents, the retrieval system entities that are responsible for the last step of the retrieval of the similar building designs. The case-based learning feature of those agents implements an own CBR domain component. This component is unique for each of the currently existing CBR agent types. It provides the corresponding agent with the reasoning functionality in order to support its decision when it comes to select the proper retrieval strategy and similarity measures.

2.5 Applying the Generic Framework Beyond Architectural Design

From the above described multi-agent-system-supported CBR-based retrieval a generic framework will be developed and applied to other domains than architecture. One specific focus will be under which constraints the generic framework can help to overcome the inherent complexity of searching for optimal subgraphs. Based on the results an according domain and task characterization will be developed. Other research focus will be dealing with the generalization of the learning agents approach for CBR-based information retrieval for design ideas generating process. The goal is to formalize and optimize the agents' experience and knowledge obtaining, teamwork and communication process in order to provide an efficient distributed case-based IR approach that is able to find information with high precision and recall rates in one or more case bases with differently (e.g. only partly) structured knowledge representation types and domain models. Consideration of applying similarity or diversity as the best suitable case comparison base will also be taken into account and a part of agents' reasoning process.

3 Current Progress

The current progress state is now in the initial phase. The research group of *Metis* is currently evaluating the user interface for creating the user queries in AGraphML (*Architectural* GraphML) format. The next steps are the integration of the interface into the retrieval system and the implementation of subgraph matching algorithms to be able to use them as second possible retrieval approach.

In the following research phase it is planned to find an explicit research direction of the PhD thesis, that can be either one of the described research foci or a combination of some of them with or without adding some new aspects that can appear during the ongoing *Metis* project discussions.

References

1. Brandes, U., Eiglsperger, M., Herman, I., Himsolt, M., Marshall, M.S.: Graphml progress report structural layer proposal. In: Graph Drawing. pp. 501–512. Springer (2002)
2. DFKI GmbH: myCBR 3 tutorial information for ICCBR 2012 (2012), http://mycbr-project.net/downloads/myCBR_3_tutorial_slides.pdf
3. KSD Research Group: KSD - Forschungsprojekte - metis (2015), http://ksd.ai.ar.tum.de/?page_id=140
4. Langenhan, C., Petzold, F.: The fingerprint of architecture-sketch-based design methods for researching building layouts through the semantic fingerprinting of floor plans. International electronic scientific-educational journal: Architecture and Modern Information Technologies 4, 13 (2010)

All links were last followed on July 8, 2015.

Interactively Learning Moral Norms via Analogy

Joseph Blass

Ph.D. Candidate, Qualitative Reasoning Group, Northwestern University, Evanston, IL
joebl@u.northwestern.edu

Abstract. Autonomous systems must consider the moral ramifications of their actions. Moral norms vary among people, posing a challenge for encoding them explicitly in a system. This paper proposes to enable autonomous agents to use analogical reasoning techniques to interactively learn an individual's morals.

1 Introduction

1.1 Challenge and Research Goals

Should a self-driving car put its passengers at risk and swerve to avoid a jaywalker, or protect its passengers and hit him? To participate in our society, computers need to share our ethics. As these systems become more autonomous, they must consider the moral ramifications of their actions. I intend to build an AI moral-reasoning system that strives for good, but can select amongst only bad options, by acquiring and applying human morals. This system will learn moral norms through natural-language interaction with humans and analogical generalization, and apply these norms by analogy.

The diversity of moral norms and concerns make hand-encoding an individual's moral sense or providing case-by-case instructions impossible. Natural interaction will be key, since users may have neither the technical skills nor understand their own morals enough to encode them themselves. Also, since human morals likely do not depend on first-principles reasoning (FPR) (Haidt, 2001), and since moral rules contradict and trade off with each other, I intend to minimize FPR in the system. A pure FPR moral reasoning system would either need rules for all possible trade-offs, to be able to ignore certain morals (a bad idea), or would freeze when moral obligations conflict. Analogical reasoning can avoid these problems if provided a good analogue.

1.2 MoralDM, Structure-Mapping, and the Companions Architecture

MoralDM (Dehghani et al. 2009) is a computer model of moral reasoning that takes in a moral dilemma in natural language, uses a natural language understanding (NLU) system to generate Research-Cyc-derived predicate-logic representations of the dilemma, and uses analogy over resolved cases and FPR over explicit moral rules to make moral decisions consistent with humans'. MoralDM is the starting point for my work.

The Structure Mapping Engine (SME), based on Gentner's (1983) Structure Mapping Theory of analogy, constructs an alignment between two relational cases and

draws inferences from it. SME can apply norms by analogy from stories (Dehghani et al. 2009). Analogy is a good fit for moral decision-making because both are guided by structure, not features. Consider the following examples. 1) A bomb will kill nine people in a room, but you can toss it outside, where it will kill one person. 2) A bomb will kill nine people, but you can toss someone onto it to absorb the blast and save the nine. Most say tossing the bomb, but not the person, is morally acceptable. These scenarios only differ structurally, in what fills which role; the entities and action types themselves are shared. The classic trolley problem (a trolley will hit five people unless it is diverted to a side track where it will hit one person), in contrast, has different features, but the same structure, as the first bomb case. Humans see these two cases as morally alike.

The Sequential Analogical Generalization Engine (SAGE) builds case generalizations that emphasize shared, and deprecate case-specific, structures. SAGE uses a case library of generalizations and exemplars. Generalizations contain facts from constituent cases: non-identical corresponding entities are replaced by abstract ones; probabilities indicate the proportion of assimilated cases each fact is present in. Given a probe, SAGE uses SME to find the most similar case in its case library. If the match is strong enough, the case is assimilated; if not, it is added as an exemplar. SAGE can use near-misses to determine defining characteristics of category members (McLure et al., 2015).

The Companion Cognitive Architecture emphasizes the ubiquity of qualitative representations and analogical reasoning in human cognition. Companion systems are designed to work alongside and interactively with humans (Forbus & Hinrichs, 2006).

2 Proposed Research and Progress

I propose to extend MoralDM in the Companion Architecture to learn to model a human user's morals. The system will learn to recognize and extract moral norms through the generalization process. It will get moral stories in natural language from the user, generate qualitative representations of those stories, generalize over those representations, and use SME to apply morals from the generalizations. I will extend MoralDM's analogical reasoning, integrate emotional appraisal, and improve NLU for a moral lexicon.

Previously MoralDM's analogical reasoning module exhaustively matched over resolved cases, which is computationally expensive and cognitively implausible. SME over ungeneralized cases also sees feature-similar but morally-different cases (i.e., the bomb scenarios) as a good match, due to the amount they have in common.

MAC/FAC is a two-step model of analogical retrieval. MAC efficiently computes dot-products between the content vectors of the probe and each case in memory (a coarse similarity measure). FAC then performs SME mappings on the most similar cases. MAC sees cases concerning mostly the same entities as the probe as good potential matches, even if the structures differ. Using MAC/FAC over generalizations rather than exemplars solves this problem, since generalizations emphasize defining structure. Abstract generalizations applied by analogy can therefore function as moral rules.

We have found that reasoning by analogy over generalizations led to more human-like judgments than using ungeneralized cases (Blass & Forbus, 2015). Reasoning can be further improved using McLure & Forbus' (2015) work on near-misses to illustrate category boundaries and the conditions for membership or exclusion. MoralDM also

still reasons using FPR about facts relevant to moral judgment, such as directness of harm. These are not explicitly stated, though we recognize them easily; MoralDM uses them in a consistency check to ensure the quality of retrieved analogues. Near-misses would let MoralDM use analogy, not FPR, to find the facts for the consistency check.

We want to expand the range and provenance of stories for MoralDM to learn from. One option is to crowd-source moral stories to present to a user for endorsement or rejection, rather than force the user to provide them all. QRG's NLU system, EA NLU, generates qualitative representations from English input, but its moral vocabulary is currently limited. The Moral Foundations Dictionary (Graham et al., 2009) is a moral lexicon; to enable EA NLU to understand moral stories, I will ensure lexical and ontological support for this vocabulary. Another NLU challenge is how to infer information implicit in the text. Work has been done at QRG on inferring narrative information, including about moral responsibility (Tomai & Forbus, 2008). I will extend EA NLU's abductive reasoning as needed to support moral narrative understanding. Finally, I will integrate emotional appraisal (Wilson et al. 2013) into MoralDM. Emotional appraisal can help recognize moral violations and enforce moral decisions.

My goal is to have a Companion running MoralDM with the above extensions interact with a human and build a model of their moral system. MoralDM could not previously do this, since it required all moral norms to be explicitly encoded, and modeled a society's aggregate judgments, not individuals. The new system will have the human tell it a moral story, crowd-source thematically similar stories, and ask the human which illustrate the same moral principle (the others are near-misses). For each story, the system would predict the moral value of actions and compare its predictions to the human's moral labels. When the core facts of the generalization stop changing and the system's labels consistently match the human's, the system has mastered that moral domain.

This project brings challenges. How much FPR will remain necessary? How must EA NLU be extended to understand moral narratives? What narrative inferences should be made about implicit information? Nonetheless, I believe I can build a system that interactively learns to model an individual's morality.

3 References

- Blass, J. & Forbus, K. (2015). Moral Decision-Making by Analogy: Generalizations vs. Exemplars. *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, Austin, TX.
- Dehghani, M., Sachdeva, S., Ekhtiari, H., Gentner, D., & Forbus, K. (2009). The role of cultural narratives in moral decision making. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society*.
- Forbus, K., & Hinrichs, T. (2006). Companion Cognitive Systems: A Step Towards Human-Level AI. *AI Magazine* 27(2), 83-95.
- Gentner, D. (1983). Structure-Mapping: A Theoretical Framework for Analogy. *Cognitive Science* 7(2).
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5), 1029.
- Haidt, J. (2001). The Emotional Dog and its Rational Tail: A Social Intuitionist Approach to Moral Judgment. *Psychological Review*, 108(4), 814-834.
- McLure, M.D., Friedman S.E. and Forbus, K.D. (2015). Extending Analogical Generalization with Near-Misses. In *Procs of the 29th AAAI Conference on Artificial Intelligence*, Austin, TX
- Tomai, E. & Forbus, K. (2008). Using Qualitative Reasoning for the Attribution of Moral Responsibility. In *Proceedings of the 30th Annual Conference of the Cognitive Science Society*. Washington, D.C.
- Wilson, J. R., Forbus, K. D., & McLure, M. D. (2013). Am I Really Scared? A Multi-phase Computational Model of Emotions. In *Proceedings of the Second Annual Conference on Advances in Cognitive Systems*.

Support to Continuous Improvement Process in manufacturing plants of multinational companies through Problem Solving methods and Case-Based Reasoning integrated within a Product Lifecycle Management infrastructure

Alvaro Camarillo

Mechanical Engineering Department, Universidad Politécnica de Madrid
a.camarillo@alumnos.upm.es

1 Introduction to problems addressed by research

The aim of this research is to capture and reuse efficiently knowledge at shop floor level of multinational companies during the resolution of manufacturing daily problems (e.g. scrap rate, quality issues, breakdowns, and in general any Continuous Improvement Process (CIP) activity). We want to provide production technicians and operators with a friendly and low time consuming Knowledge Management (KM) tool to get their engagement and collaboration, avoiding negative impact in productivity, and promoting the knowledge share across plants overcoming language, nationalities, and competition barriers.

We propose the Problem Solving (PS) method 8D as structured process to guide the knowledge share. A Product Lifecycle Management (PLM) system will be the logical infrastructure to store all product, process, machinery, and users information. This PLM system will host also the database of a Case-Based Reasoning (CBR) system. This CBR system will be the KM tool in charge of capturing and reusing the knowledge [1],[4],[5],[6]. FMEAs of Design, Process and Machinery will be used to populate initially the CBR System [3].

The CBR cycle [1] would be though as follow:

- User introduces basic description of de problem (new case).
- Based on this description the CBR system collects additional information related product, process, machinery or users from the PLM. It proceeds to calculate to find similar cases.
- The system proposes containment actions and different root causes (retrieved cases). The user checks these root causes in the line and gives feedback to system.
- Based on the corrected list of most similar cases the system performs adaptation and proposes a solution (solved case).
- Solution is tested by user (tested/repared case) and implemented together with its associated preventive actions.

- The learned case is stored in the database of cases.

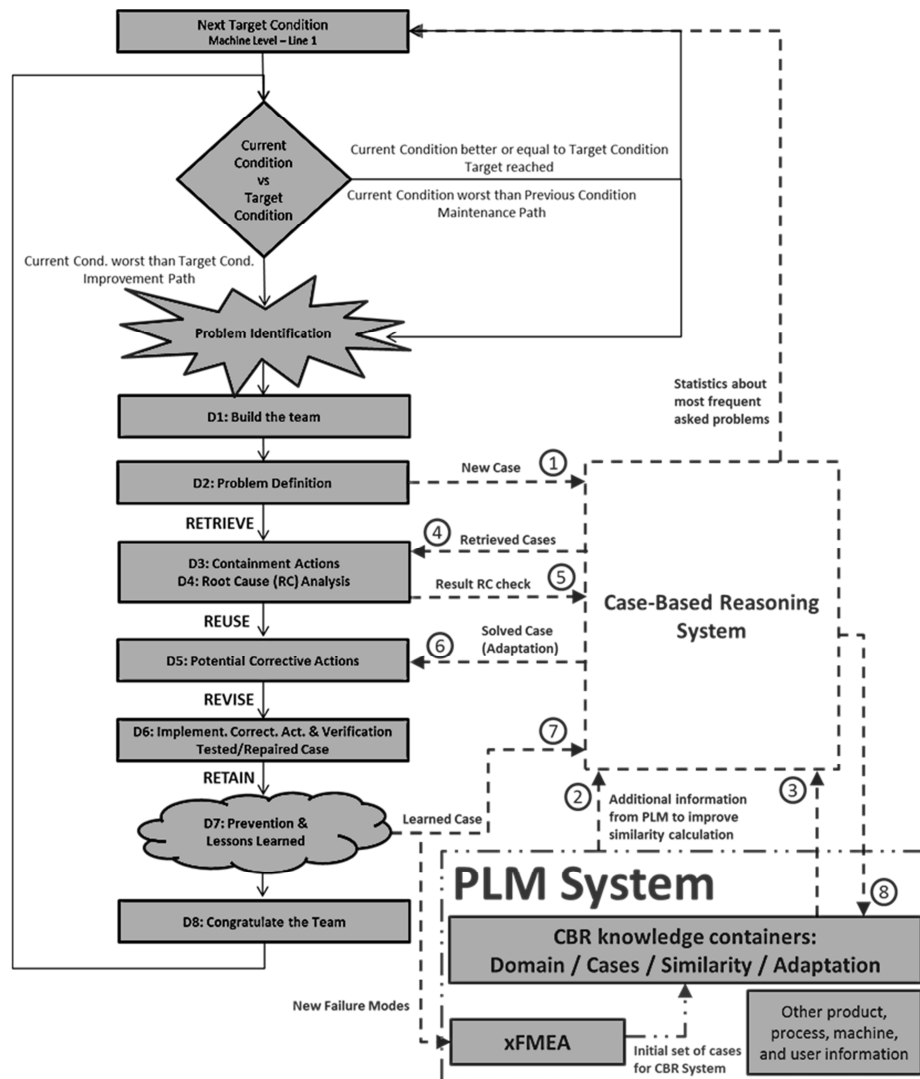


Fig. 1. – Proposed model for supporting CIP through PS, CBR and PLM

Expected contributions of this research:

- Combination in a single model of PS methods, as process for guiding the share of knowledge, CBR, as KM tool, PLM as global infrastructure to contain information and control information flow, and FMEA method, as the tool to define manufacturing problems in a formalized way and to populate initially the CRB System.

- Based on minimal data introduced by the user (low time consumed) we propose to get from the PLM extended information that will be used to calculate similarity.
- Bring this type of KM tool not only to designers or to engineers, but direct to blue-collar associates working at production lines.

2 Description of progress to date

Currently we are developing the Model that has to support the knowledge capture and reuse (see fig. 1). For the case study, two open source applications have been selected: Aras as PLM software (www.aras.com), and myCBR as CBR software (www.mycbr-project.net) [2]. For the implementation, a multinational company of the electrical batteries branch was selected. To get the benefits of knowledge sharing between two teams with very low interaction until now the system will be installed in two manufacturing plants located in two different countries. It will focus only on one of the production steps of batteries in order to get consistent results in a limited period of time.

3 Proposed plan for research

After the review of the state of the art in the fields of CIP, PS, CBR, and PLM, we are currently designing the initial knowledge containers of domain, similarity and adaptation of the CBR system [6] that will be used to test our concept initially in a single production line. This task has to be finished by the end of May 2016. The experience from this initial test will be used to improve our KM tool in order to do a second test loop at whole plant level. Finally a third test loop will be performed between the two plants until end of 2016. The presentation of the PhD Thesis is planned in May 2017.

References

1. Aamodt, A., Plaza, E.: Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. In: *AI Communications*, vol. 7, no. 1, pp. 39-59 (1994)
2. Atanassov, A., Antonov, L.: Comparative Analysis of Case Based Reasoning Software Frameworks jCOLIBRI and myCBR. In: *Journal of the University of Chemical Technology and Metallurgy*, vol. 47, no. 1, pp. 83-90 (2012)
3. Atamer, A.: Comparison of FMEA and field-experience for a turbofan engine with application to case based reasoning. In: *Aerospace Conference, 2004. Proceedings. 2004 IEEE*, Vol. 5, pp. 3354-3360 (2014)
4. Jabrouni, H., Kamsu-Foguem, B., Geneste, L., Vaysse, C.: Continuous Improvement Through Knowledge-Guided Analysis in Experience Feedback. In: *Engineering Applications of Artificial Intelligence*, vol. 24, no. 8, pp. 1419-1431 (2011)
5. Kamsu-Foguem, B., Couderta, T., Bélera, C., Geneste, L.: Knowledge Formalization in Experience Feedback Processes: An Ontology-Based Approach. In: *Computers in Industry*, vol. 59, no. 7, pp. 694-710 (2008)
6. Richter, M., Weber, R.: *Case-Based Reasoning: A Textbook*. Springer, Heidelberg (2013)

Aspect-based Sentiment Analysis for Social Recommender Systems

Yoke Yie Chen

IDEAS Research Institute,
Robert Gordon University,
Aberdeen, Scotland
`{y.y.chen}@rgu.ac.uk`

1 Introduction

Social recommender systems provide users with a list of recommended items by exploiting knowledge from social content. Representation, similarity and ranking algorithms from the Case-Based Reasoning (CBR) community have naturally made a significant contribution to social recommender systems research [1, 2]. Recent works in social recommender systems have been focused on learning implicit preferences of users from online consumer reviews. Most online reviews contain user opinion in the form of positive and negative sentiment on multiple aspects of the product. Since a product may have multiple aspects, we hypothesize that users purchase choices are based on comparison of products; which implicitly or explicitly involves comparison of aspects of these products. Therefore, our main research question is “*Does considering product aspects importance (weight) improve prediction accuracy of a product ranking algorithm?*”

2 Research Aim

This research aims to develop a novel aspect-based sentiment scoring algorithm for social recommender systems. Our particular focus will be on using social content to develop novel algorithms for different product domains. For this purpose, we intend to:

1. Develop an aspect extraction algorithm to extract product aspects.
2. Develop aspect weighting algorithms to extract product aspects weights from social content.
3. Study the effect of temporal dynamics on aspect weight.
4. Evaluate the performance of our proposed algorithms in performing a top-N recommendation task using standard performance metrics such as mean average precision.

3 Challenges

Social recommender system harness knowledge from product reviews to generate better recommendation. Key to this task is the need for a novel aspect based sentiment analysis approach to harness this large volume of information. However, this approach suffers three main challenges:

1. Aspects extracted from product reviews using NLP-based techniques rely on POS tagging and syntactic parsing which are known to be less robust when applied to informal text. As a result, it is not unusual to have a large numbers of spurious content to be extracted incorrectly as aspects.
2. A user's purchase decision hints at the aspects that are likely to have influenced their decision and as such be deemed more important. To understand the importance of an aspect to users, it is necessary to further reveal the importance weight that users placed on an aspect. Additionally, user preferences change over time. Term frequency (TF) is the naive approach for this task where the weight of an aspect is equal to the number of occurrences of that aspect in product reviews. However, this approach is not able to capture users' preferences that change over time.
3. The absence of ground truth data causes evaluating ranking algorithm a challenging task in recommender system. For example, Best Seller ranking in Amazon can be a straightforward reference to evaluate the ranking of system generated recommendation list. However, this ranking is biased towards old products in Amazon. Therefore, there is a need to study relevant knowledge sources to construct a reference ranking for evaluation purpose.

4 Proposed Plan of Research

To answer our research question, our proposed plan of research is:

1. Compare the performance of our proposed aspect extraction algorithm with key state-of-the-art algorithms to determine the impact of aspect quality on recommendation tasks. We will evaluate the performance of these algorithms through accuracy metrics in extracting genuine product aspects. Thereafter, we apply the extracted aspects from these approaches in our aspect based sentiment scoring algorithm and rank the products. We then compare the recommendation performance of aspect based sentiment scoring algorithm with a sentiment analysis algorithm that is agnostic of aspects.
2. Feature selection techniques in machine learning are known to enhance accuracy in supervised learning tasks such as text classification by identifying redundant and irrelevant features. We propose to explore different feature selection techniques (e.g. Information Gain and Chi-squared) to select aspects that are important to users.
3. Our initial approach in aspect weighting algorithm places individual product aspect with equal importance weight across all products. We intend to

explore other related approaches such as TF-IDF (Term Frequency Inverse Document Frequency) to represent the importance of a product aspect. TF-IDF has been widely used in Information Retrieval community to evaluate the importance of a word to a document in a corpus. We propose to augment our aspect weighting algorithm by evaluating the importance of a product aspect to a particular product.

4. To study the effect of temporal dynamics in aspect importance weight, we look into investigating aspect weights that are inferred by:
 - **Trending information.** We would like to analyse different trending patterns of aspects occurrence in product reviews over the years (e.g. upward, downward and recurring trend). Specifically, a higher weight should be given to aspects which have an upward and recurring trend, indicating that the importance of an aspect is growing. Likewise, a lower weight should be given to aspects having a downward trend.
 - **Recency of aspects.** Aspects which frequently appear in old product reviews will have a lower weight than aspects appearing in recent product reviews. This indicates that aspects that are frequently occurring in recent product reviews are deemed important.
5. To evaluate our ranking algorithm, we use users' ratings as the baseline to compare with our proposed ranking approach. This baseline ranks each product using the average users' rating. Products in the higher rank are thus recommended.

5 Current Progress

Designed and developed novel algorithms in the following areas:

- **Aspect extraction.** The proposed approach integrates semantic relationship and frequency cut-off. The proposed approach was evaluated against state-of-the-art techniques and obtained positive results.
- **Aspect selection.** We address the problem of selecting important aspects using feature selection heuristics based on frequency counts and Information Gain (IG) to rank and select the most useful aspects.
- **Aspect-based sentiment scoring.** The proposed algorithm incorporates aspect importance weight and sentiment distribution. We investigated two different resources that infer the importance of product aspects: preference and time.

References

1. R. Dong, M. Schaal, M. OMahony, K. McCarthy, and B. Smyth. Opinionated product recommendation. In *Inter. Conf. on Case-Based Reasoning*. 2013.
2. L. Quijano-Sánchez, D. Bridge, B. Díaz-Agudo, and J. Recio-García. Case-based aggregation of preferences for group recommenders. In *Case-Based Reasoning Research and Development*, pages 327–341. 2012.

Toward a Case-Based Framework for Imitation Learning in Robotic Agents

Tesca Fitzgerald

School of Interactive Computing,
Georgia Institute of Technology,
Atlanta, Georgia, 30332
`tesca.fitzgerald@cc.gatech.edu`

1 Introduction

Imitation learning is a skill essential to human development and cognition [6, 5]. Naturally, imitation learning has become a topic of focus for robotics research as well, particularly in interactive robots [1, 2]. In imitating the actions of a teacher, a cognitive agent learns the demonstrated action such that it may perform a similar action later and achieve a similar goal. Thus, we expect that a cognitive robot that learns from imitation would reuse what it has learned from one experience to reason about addressing related, but different, problem scenarios.

The eventual goal of this work is to use a case-based approach to enable imitation learning in interactions such as the following. A human teacher guides the robot to complete a task, such as scooping the contents of one container into another. The robot records the demonstrated actions and observed objects, saving the demonstration as a *source case* in its case memory. At a later time, the robot is asked to repeat the *scooping* task, but in a new, *target* environment containing a different set of object features to parameterize and execute the task. Next, the robot would transfer its representation of the *scooping* task to accommodate for the differences between the source and target environments, and then execute an action based on the transferred representation to achieve the goal state in the target environment.

Using a case-based framework to address this problem allows us to represent demonstrations as individual experiences in the robot's case memory, and provides us with a framework for identifying, transferring, and executing a relevant source case demonstration in an unfamiliar target environment. The main research questions we plan to address are as follows:

- How should task demonstrations be represented in case memory?
- How do we determine which features of a robot's environment are relevant to completing a task, and thus should be stored in the source case?
- What features should be considered in retrieving a source case demonstration for reuse in a target environment? How should these features be prioritized during source case retrieval?

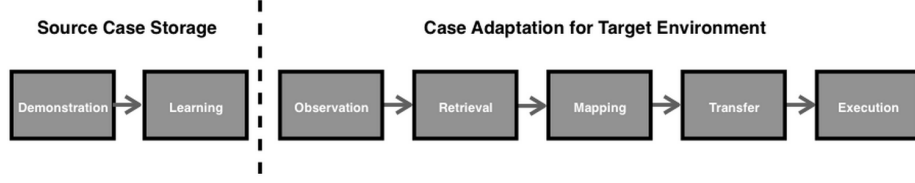


Fig. 1. Case-Based Process for Task Demonstration Transfer

2 Research Plan and Progress

We have defined a case-based approach to transfer for enabling imitation in robotic agents, consisting of two separate processes (as shown in Figure 1): the *Case Storage* process in which the robot receives demonstrations of a task and stores each demonstration as a case in source memory, and a *Case Adaptation* process which is used at a later time when the robot is asked to repeat a task in a target environment.

2.1 Case Storage Process

Demonstration and Learning We have implemented the first step in the *Case Storage* process, where the robot records and stores each task demonstration as a source case in memory. We define each case as the tuple $C = \langle L, D, T, O, S_i, S_f \rangle$, where:

- L represents the label of the task which was demonstrated, e.g. "scooping".
- D represents the set of action models which encode the demonstrated motion, represented as Dynamic Movement Primitives as defined in [4].
- T is the set of parameterization functions which relate the set of action models to the locations of objects in the robot's environment. For example, a parameterization function may be used to represent how the robot's hand must be located above a bowl prior to completing a *pouring* action.
- O is the set of *salient* object IDs which are relevant to the task.
- S_i and S_f are the initial and final states, respectively, which represent the set of objects observed in an overhead view of the robot's environment.

2.2 Case Adaptation Process

At a later time, the robot may be asked to repeat a learned task in an unfamiliar target environment. Using the framework shown in Figure 1, the robot may address a target environment using the following steps.

Observation The robot is given a target problem to address, under the assumption that it has a relevant source case in memory which can be used to address the target problem. The robot observes the target environment by viewing the objects located in the table-top environment using an overhead camera, providing it with the initial state S_i of the target case.

Retrieval and Mapping The robot must then choose a source case from memory containing the demonstration that is most relevant to the current target problem. Once a relevant source case has been retrieved, a mapping must be generated that encodes the differences between the source and target environments. This mapping is later used to transfer the source case such that differences in the target environment are addressed. We have not yet implemented the Retrieval and Mapping steps, but will be addressing them in upcoming work.

Transfer and Execution We have implemented the last two steps of the Case Adaptation process, the *Transfer* and *Execution* steps. Currently, we manually provide the robot with the most relevant source case demonstration and a mapping between objects in the source and target environments.

We take a similarity-based approach to transfer, where we consider the similarity between the source case and target environments when defining transfer processes. As we encounter transfer problems in which the source and target problems become less similar, the source case is transferred at a different level of abstraction, such that only high-level features of that case are transferred. We have implemented three transfer methods, each of which operates by transferring the source case at a different level of abstraction (further described in [3]). Once the source case has been transferred, it is used to plan and execute a new action trajectory to address the target problem. Preliminary experiments have evaluated each method under the assumption that we select the approach, and thus the level of abstraction at which transfer occurs, to be used for a given transfer problem.

3 Future Work

Our current implementation assumes that we manually provide a mapping between equivalent objects in the source and target environments. We plan to identify (i) a method for autonomously determining this object mapping and (ii) a process for identifying and retrieving an appropriate source case demonstration.

References

1. Argall, B.D., Chernova, S., Veloso, M., Browning, B.: A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57(5), 469–483 (2009)
2. Chernova, S., Thomaz, A.L.: Robot learning from human teachers. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 8(3), 1–121 (2014)
3. Fitzgerald, T., Goel, A.K., Thomaz, A.L.: A similarity-based approach to skill transfer. *Women in Robotics Workshop at Robotics: Science and Systems* (2015)
4. Pastor, P., Hoffmann, H., Asfour, T., Schaal, S.: Learning and generalization of motor skills by learning from demonstration. In: *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. pp. 763–768. IEEE (2009)
5. Piaget, J., Cook, M.T.: The origins of intelligence in children. (1952)
6. Tomasello, M., Kruger, A.C., Ratner, H.H.: Cultural learning. *Behavioral and brain sciences* 16(03), 495–511 (1993)

Doctoral Consortium Research Summary: Virtuosity in Computational Performance

Callum Goddard

School of Electronic Engineering and Computer Science,
Queen Mary University of London, Mile End Road,
London E1 4NS, United Kingdom
c.goddard@qmul.ac.uk
<http://www.eecs.qmul.ac.uk/>

Abstract. This is a research summary of Virtuosity in Computational Performance, addressing the question: *How can a computer, as judged by a human audience, demonstrate virtuosity in computational performances with a physical model of a bass guitar?* The proposed plan for this research is to develop a computational performance system which uses case-based reasoning and reflection to produce virtuosic performances with a physical model of an electric bass guitar. Three supporting studies are planned to investigate bass playing, collect performance data and perceptions of virtuosity.

Keywords: Computational Performance, Virtuosity, Case-based Reasoning, Reflection, Physical Modelling, Music Analysis

1 The Problem being Addressed and Research Questions

The main question this research is addressing is:

How can a computer, as judged by a human audience, demonstrate virtuosity in computational performances with a physical model of a bass guitar?

Computationally performed music, where a computer is responsible for rendering, generating or synthesising the music in its entirety, can appear lacking, robotic or sterile [1]. There are approaches to overcome this that focus on introducing or emulating expressive phrasing within a performance of a score [1]. However, if instead of expression virtuosity was exhibited within a computer performance, would this not offer a more satisfying solution to sterile performances as well as aiding in investigations into virtuosity of human performances?

Virtuosity here is being viewed as a property of a performance, formed through a complex and dynamic relationship between the performer, an audience and the domain in which the performance is situated [2]. It encompasses notions of expression and style within the performance alongside a demonstration of high levels of technical proficiency, a deeper understanding of the instrument, the piece being played and the context or domain of the performance.

The decision to limit the scope of this research to the domain of electric bass has been made as the author is an experienced electric bass player. There is also recent research [3, 4] within this area that can be used within this PhD.

2 Proposed Plan for Research

To address the main research question, this research will focus on developing a theory for how a computer can exhibit virtuosity within a rendered performance. To allow this theory to be tested a computer performance system that can create performances, using the physical model of electric bass guitar developed by Kramer et al. [3], is planned to be developed.

The current theory is based upon a case-based reasoning approach. Previous work on the SaxEx system [5] has demonstrated how effective case-based reasoning can be when applied to creating expressive performances. Unlike the SaxEx system, which manipulates the waveforms of a non-expressive audio recording as its output, the planned system will be manipulating physical model parameters. These parameters are intended to be abstracted to allow for rationalisation of the performances and evaluation of their virtuosity.

A performance here is being formalised in Equation 1 as the result of *Player* applying a set of *Techniques*, $\{T_{pluck}, T_{thumb} \dots T_n\}$, to a sequence of *Notes*, $\langle N_1, N_2 \dots N_n \rangle$. Musical score information, physical model and performance parameters are needed to be represented, abstracted and manipulated to produce a performance. All this information will be represented using the Common Hierarchical Abstract Representation for Music (CHARM) [7, 8].

$$Performance = Player(Techniques, Notes) \quad (1)$$

Cases are to be CHARM constituents. Constituents are formed by grouping together particles. Particles can be either events and/or other constituents. An event differs from a constituent in that it is the most fundamental element of interest within the data and as such cannot be formed from groupings. Constituents enable the formation of hierarchical structures, denoting increases in both hierarchical level and in abstraction. Events form the lowest levels of this hierarchy and within this research will be musical notes. A visual representation of an example is constituent is shown in Figure 1.

When producing a new performance, or interpreting one, a new CHARM representation will need to be constructed. First, constituents of suitable types are found, or created, and then searches for similar constituents are made. A constituent's similarity is to be judged on its structural and musical type, along with the combination and type of its particles. Retrieved constituents can be modified by interchanging particles for better matching ones to increase the suitability of the constituent for the new case. This process of finding new constituents, then modifying them is akin to the engagement reflection cycle outlined by Pérez y Pérez [6], and is important as being able to reflect upon the performances the system creates can help to guide it towards producing virtuosic one.

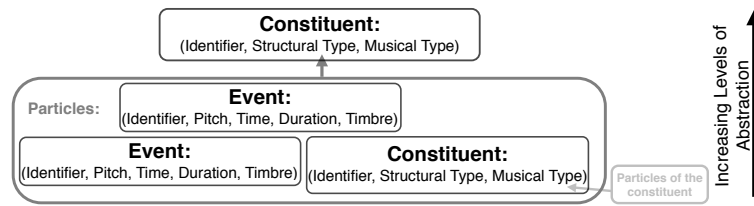


Fig. 1. A visual representation of a CHARM constituent formed of a group of three particles: two events and a constituent, (which has its own particles). I refer the reader to Smaill et al. [7] for more details on the internal structure of event and constituents.

Ontologies for domain specific knowledge e.g. musical score structure, bass technique etc. will be separate from the CHARM representation forming add-on modules for the system. To further inform the knowledge required by the system three studies are planned. One to investigate aspects of bass playing, one to collect performance data and third to see how virtuosity is perceived to inform the reflection of the system.

3 Description of Progress to date

At present I am approaching the first year review of my PhD. The work so far has been in better understanding the form the PhD will be taking, with this document forming a brief summary of the work that has been completed so far.

Acknowledgments. This work is supported by the Media and Arts Technology programme, EPSRC Doctoral Training Centre EP/G03723X/1

References

1. Widmer, G. and Goebel, W. Computational models of expressive music performance: The state of the art. *JNMR*, 33(3):203-216. (2004)
2. Howard, V. A.: *Charm and speed: virtuosity in the performing arts*. Peter Lang Publishing Inc., New York. (2008)
3. Kramer, P., Abesser, J., Dittmar, C., and Schuller, G.: A digitalwaveguide model of the electric bass guitar including different playing techniques. *ICASSP, IEEE* (2012)
4. Abeßer, J.: Automatic Transcription of Bass Guitar Tracks applied for Music Genre Classification and Sound Synthesis. PhD thesis, Elektrotechnik und Informationstechnik der Technischen Universität Ilmenau (2014)
5. Arcos, J. L., Mántaras, R. L., and Serra, X.: Saxex: a case-based reasoning system for generating expressive musical performances. *JNMR*, 27(3):194-210. (1998)
6. Pérez y Pérez, R.,. *MEXICA: A Computer Model of Creativity in Writing*. PhD thesis, The University Of Sussex. (1999)
7. Smaill, A., Wiggins, G., and Harris, M.: Hierarchical music representation for composition and analysis. *CHum*. (1993)
8. Wiggins, G., Miranda, E., Smaill, A., and Harris, M.: A framework for the evaluation of music representation systems. *CMJ*, 17(3):31-42. (1993)

System to design context-aware social recommender systems

Jose L. Jorro-Aragoneses

Department of Software Engineering and Artificial Intelligence
 Universidad Complutense de Madrid, Spain
 email: jljorro@ucm.es, belend@ucm.es, jareciog@fdi.ucm.es

Abstract. In this document, we summarize my PhD thesis goals and the progression in 2014/2015. The principal goal of my PhD thesis is to describe an architecture to design context social recommender systems. Finally, we explain all goals that we will try to achieve during my PhD studies.

1 Introduction

The number of products and the amount of information that we can consider has increased with the growth of the Internet. Sometimes, all of this information could overwhelm users. Recommender systems were created to filter this information and they just show the most interesting results for each user. For example, recommender systems are an important feature in e-commerce, where they show what products may most interest a user [8].

Recommender systems are an active research area in the artificial intelligence community. The majority of recommender systems use features of products and user preferences to calculate recommendations [7]. A trend in this area is to use contextual information [1] in recommender systems.

We find a complete definition of context in [3] “*Context is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves*”. In our case, entities of recommender systems are items that systems recommend and users who receive a recommendation.

My PhD thesis goal is to study what kinds of context information there are in a recommender system, how many ways we can obtain this information (implicit, users introduce the information, or explicit, the system obtains this information itself) and design a system to build recommender systems automatically.

The paper is organized as follows: Section 2 defines specific objectives in my PhD thesis based on the main goal. Finally, we explain the progress to date in Section 3.

2 Research objectives

As we said before, the main goal of my PhD thesis is to analyse what type of context information could be used in a recommender system and we will use these results to create an architecture that creates templates of CBR recommender systems automatically. To do it, we need to analyse different recommender systems and observe what type of information have its elements.

We can find 4 types of context information based on [3]:

Individual: Features of entities (age, sex, restaurant type, ...).

Location: Location of entities (longitude, latitude, room of a museum, ...).

Time: Time or time restrictions of entities (timetable, date of an event, ...).

Relationship: Features that we obtain in entities relationship (a family, a group of pictures of the same theme,...).

Currently, we are studying recommender systems that we have built before and classifying context information of items and types and type of users. Firstly, we classify *MadridLive* [2, 6], a recommender system of tourism and leisure activities in Madrid. This system uses all types of context information, and after, we add the emotional context (part of my PhD thesis) to complete the system. At the same time, we study different ways to obtain this information (mobile devices, social networks, linked-data, etc.). With context types and forms to obtain the information we create an ontology that classifies CBR recommender systems by the type of information that these systems use. Finally, we are going to use this ontology to make a system that builds templates of CBR systems. This system will use the type of items, users and technology to create a template that explains how to build the recommender system.

Preliminary specific objectives are defined as follows:

Objective 1: Detection and study of the influence of emotional context in recommender systems.

Objective 1.1: Obtain a method to detect the user emotions by his/her facial gestures. *The preliminary results have been published in [5, 4].*

Objective 1-2: Investigate different applications of emotions in recommender systems.

Objective 2: Classify all context information types and all different forms to obtain each type of information. To do it, we will design an ontology that will be used to create the final system. In this objective we are going to study recommender systems for individuals only.

Objective 3: Extend classification to group recommender systems. The main goal is to detect and study the social context in recommender systems.

Objective 3.1: Obtain a method to calculate the influence between members of a group using social networks.

Objective 3.2: Determine if there are patterns of groups with similar characteristics, for example, families, seniors group, etc.

Objective 3.3: Add social context conclusions in the ontology that we have defined in **Objective 2**.

Objective 4: Design a system that uses our ontology to create templates for CBR systems. The system creates templates using the information that the recommender system will use.

Objective 5: Make experiments to validate the system and research the influence of each type of context in a recommender system. To do it, we are going to create a recommender system based in the tourism and leisure domain.

All these specific objectives permit us to study all information types that participate in a recommender system and propose a system to design recommender systems automatically.

3 Description of the progress to date

In 2014/2015, I have finished objective-1.1. I have proposed a CBR approach to infer the emotion state using images of the user's face. This method has been published in [5] and [4]. Next, we have compared the quality of our method with others, and this comparison is explained in the paper that we have published at this conference (ICBBR 2015). In this paper, we explain a possible solution to the cold-start problem. To do it, we have created specialized case bases with cases that have the same features. These features are:

- *Age*, classified in two categories, children and adults.
- *Gender*, classified in two categories, men and women.
- *Ethnic group*, classified in the ethnic group features as Japanese, European, etc.

Actually, I am studying the design of an ontology to classify recommender systems by the type of context information that they use. The objective is using the ontology in a system that creates templates of CBR systems.

References

1. Adomavicius, G., Tuzhilin, A.: Context-Aware Recommender Systems, chap. 7, pp. 217–253. Springer US (2011)
2. Agudo Calvo, B., Jorro Aragoneses, J.L., Valle Corral, J.d.: Sistema de recomendación de actividades turísticas: Madrid live (2013)
3. Dey, A.K.: Understanding and Using Context. *Personal and Ubiquitous Computing* 5(1), 4–7 (Feb 2001), <http://dl.acm.org/citation.cfm?id=593570.593572>
4. Jorro-Aragoneses, J.L., Díaz-Agudo, B., Recio-García, J.A.: Optimization of a CBR system for emotional tagging of facial expressions. In: UKCBR (2014)
5. López-de-Arenosa, P., Díaz-Agudo, B., Recio-García, J.A.: CBR tagging of emotions from facial expressions. In: ICCBR (2014)
6. López Gómez, J., Bezares Álvarez, M., Marín Fernández, R.: Madrid live: un sistema de recomendación de ocio social y contextual para la ciudad de madrid (2014)
7. Ricci, F., Rokach, L., Shapira, B.: Introduction to Recommender Systems, chap. 1, pp. 1–35. Springer (2011)
8. Schafer, J., Konstan, J., Riedl, J.: E-commerce recommendation applications. In: Kohavi, R., Provost, F. (eds.) *Applications of Data Mining to Electronic Commerce*, pp. 115–153. Springer US (2001)

Experience-based Recommendation for a Personalised E-learning System

Blessing Mbipom

School of Computing Science and Digital Media,
Robert Gordon University,
Aberdeen, Scotland,
United Kingdom.
`b.e.mbipom@rgu.ac.uk`

1 Introduction

A large amount of learning resources is available to learners on the Web. Users of these resources are often discouraged by the time spent in finding and assembling relevant resources to support their learning goals, and these users often face the information overload problem [4]. Personalisation within e-learning would allow the learning abilities and preferences of individual learners to be taken into account, thus enabling such systems to offer relevant resources to learners [7].

The interaction of previous learners with resources and the resulting outcome can be viewed as a learning experience. An experience-based recommendation approach would allow the experiences of similar users to be reused for making recommendations to new users. Currently, some recommender systems in e-commerce can capture the experience of users with items and reuse these to enhance recommendation [2]. However, little work has been done to reuse experiences in the e-learning domain [1, 3]. There is potential to improve the recommendations made within e-learning [5], drawing from the impact that the reuse of user experiences has made within e-commerce. Although recommendation in the e-learning domain is challenging given that the learning resources have to be carefully combined unlike individual products in e-commerce.

A key contribution of this research will be the development of innovative approaches to incorporate the learning experiences of previous learners captured in outcomes such as reviews and ratings, in the recommendations made to new learners. This research will harness the wide range of available e-learning resources in order to cater for learners with different preferences. The knowledge contained in the learning resources will be employed for refining learners' goals and indexing new learning resources. This work will improve the current state of e-learning systems by reusing the experiences of previous learners when recommending relevant learning resources to new learners.

2 Research Questions

This research aims to capture and reuse the learning experiences of previous learners to enhance recommendations made to new learners within a personalised e-learning system? This research seeks to address the following questions:

- How can learners' goals be refined to improve the recommendation of learning resources?
- How can learners' preferences and abilities be captured to enhance personalised recommendations?
- How can learning resources be represented to support effective retrieval?
- How can outcomes such as learners' reviews and ratings, be captured and reused to enhance e-learning recommendation?

3 Research Plan

This research will involve the development of novel approaches for reusing the experiences of previous learners to enhance e-learning recommendation. Techniques to capture learners' preferences and abilities will be developed. Existing learner models will be adapted for this task with the aim of capturing the preferences and the abilities of learners. This information would be used for making relevant recommendations to new learners.

Existing knowledge sources will be organised into a coherent background knowledge structure. Potential knowledge sources such as Microsoft Academic Search, the ACM Computing Classification System, and Wikipedia have already been identified. The plan is to employ these in the development of a background knowledge structure which can be employed for refining learners' goals and for indexing learning resources. This structure will be useful for identifying the links between resources and for recommending relevant resources.

Methods for representing and refining learners' goals will be developed. This is necessary in e-learning because learners often have insufficient knowledge of the domain to formulate suitable goals. The plan is to map the goals to a resource representation developed using shared background knowledge, this will entail reasoning with the text in the goals and the learning resources.

Representations that capture learners' outcomes will be created. Learners' test scores, reviews and ratings can be viewed as outcomes in an e-learning domain. Currently, learners' test scores are the major form of feedback used in e-learning. However, this does not capture learners' opinions which can be effectively employed to inform other learners. The plan is to incorporate user-opinions with user-performance to enhance the recommendation process.

4 Current Progress

The research methodology has been substantially developed. Various approaches for representing learning resources have been investigated, these range from

knowledge-light to knowledge-rich approaches. Some methods of refining learners' goals have also been examined.

Different types of learning resources have been identified to use as data for this work. These include e-books, online teaching slides and video lectures. They have been chosen because they contain structure and metadata that will help with the research, and because of the variety of media types contained.

Preliminary experiments have been carried out to develop a background knowledge structure to use for the refinement of learners' goals and the representation of new learning resources. A collection of 217 e-book chapters from the machine learning domain were collected for the experiments. Terms and phrases were extracted from the Tables-of-contents (TOCs) of the e-books using some NLP techniques and phrase identification methods.

E-books are used as the primary data source in this work because of the structure they contain and because they are designed to be effective for teaching and learning. Furthermore, issues of trust and provenance [6] are catered for because the nature of books means an author and affiliation exists. Wikipedia is used as a complementary data source, because it is a knowledge-rich source put together by many contributors.

Terms and phrases were extracted from the TOCs of e-books and compared with phrases from the Machine Learning category in Wikipedia to generate a set of suitable phrases to use for developing the background knowledge structure. The result was 90 phrases consisting of 17 unigrams, 58 bigrams and 15 trigrams.

Initial output shows the potential to harness the knowledge in e-Books and Wikipedia for developing a background knowledge structure that will enable the refinement of learners' goals and indexing of new learning resources. Further work will involve evaluation of this method, and the development of a system that employs the background structure to recommend relevant learning resources.

References

1. Bobadilla, J., Hernando, A., Arroyo, A.: E-learning experience using recommender systems. In: Proceedings of the 42nd ACM Technical Symposium on Computer Science Education, pp. 477–482. ACM (2011)
2. Dong, R., Schaal, M., O'Mahony, M.P., McCarthy, K., Smyth, B.: Opinionated product recommendation. In: S.J. Delany, S. Ontañón (eds.) Case-Based Reasoning Research and Development, *LNC3*, vol. 7969, pp. 44–58. Springer, Heidelberg (2013)
3. Ghauth, K.I.B., Abdullah, N.A.: Building an e-learning recommender system using vector space model and good learners average rating. In: Ninth IEEE International Conference on Advanced Learning Technologies, pp. 194–196. IEEE (2009)
4. Kantor, P.B., Rokach, L., Ricci, F., Shapira, B.: Recommender systems handbook. Springer (2011)
5. Kolodner, J.L., Cox, M.T., González-Calero, P.A.: Case-based reasoning-inspired approaches to education. *The Knowledge Engineering Review* **20**(3), 299–303 (2005)
6. Leake, D., Whitehead, M.: Case provenance: The value of remembering case sources. In: Case-Based Reasoning Research and Development, pp. 194–208. Springer (2007)
7. Peter, S.E.: The use of tagging to support the authoring of personalisable learning content. Ph.D. thesis, University of Greenwich (2012)

Workflow Adaptation in Process-oriented Case-based Reasoning

Gilbert Müller

Business Information Systems II
University of Trier
54286 Trier, Germany
muellerg@uni-trier.de,
<http://www.wi2.uni-trier.de>

Workflows are an important research domain, as they are used in many application areas, e.g., there are business workflows, scientific workflows, workflows representing information gathering processes, or cooking instructions. Workflows are “the automation of a business process, in whole or part, during which documents, information or tasks are passed from one participant to another for action, according to a set of procedural rules” [4]. Thus, workflows consists of a structured set of tasks and data objects shared between those tasks. In this regard, Process-oriented Case-based Reasoning (POCBR) [7] addresses the creation and adaptation of processes that are, e.g., represented as workflows. Although, POCBR is of high relevance little research exist so far.

The presented research focuses on the development of new workflow adaptation approaches and related topics, for instance the retrieval of workflows. Methods are investigated, which automatically learn adaptation knowledge from the case base. This prevents limited adaptation capabilities due to the acquisition bottleneck for adaptation knowledge.

1 Research Questions

This section presents the research questions addressed by my doctoral thesis in note form.

1. How can workflows be efficiently retrieved?
2. How can workflows be adapted regarding defined preferences or restrictions?
3. How can interactive workflow adaptation be realized?
4. How can the adaptability of workflows be reflected during retrieval?
5. How can adaptation knowledge be revised to address the retainment of adaptation knowledge?

The approaches to address the first two research questions are described in the next section and section 3 describes how the remaining open research questions are going to be investigated.

2 Current state of research

The presented research is implemented and evaluated using the CAKE (Collaborative Agent-based Knowledge Engine) framework¹ developed at the University of Trier. It deals with semantic workflows and is able to compute the similarity between two workflows according to the semantic similarity [2]. The approaches will be illustrated and investigated in the cooking domain, i.e., the workflows represent cooking recipes.

Currently, approaches addressing the first two research questions have been investigated:

1. Based on research about clustering of workflows [3], the problem of improving retrieval performance by developing a cluster-based retrieval method for workflows [8] was addressed. To achieve this, a new clustering algorithm, which constructs a binary tree of clusters was developed. The binary tree is used as index structure during a heuristic search to identify the most similar clusters containing the most similar workflows in a top-down fashion. Further, POQL [12] was developed serving as query language to guide the retrieval and the adaptation of workflows regarding defined preferences or restrictions.
2. Several adaptation approaches had been investigated to address the second research question. A compositional adaptation approach for workflows was investigated [9] where workflows are decomposed into meaningful subcomponents, called *workflow streams*. In order to support adaptation, streams of the retrieved workflow are replaced by appropriate streams of other workflows. Based on this work, operator-based adaptation [11] has been developed. The adaptation operators are learned automatically based on the workflows in the case base enabling to remove, insert or replace workflow fragments. Further, workflow generalization and specialization has been addressed [10], which increases the coverage of the workflow cases and thus being able to support adaptation as well.

3 Future Work

In future work, an additional adaptation approach will be investigated for semantic workflows, similar to the adaptation approach presented by Minor et. al. [6], which is based on adaptation cases describing how to transform a particular workflow to a target workflow. The future work addressing the remaining research questions 3.-5. is summarized below.

A drawback of applying traditional adaptation methods is that the adaptation goal must mostly be known previously. Consequently, this can lead to a non-optimal or not desired solution. Hence, interactive adaption [1] will be investigated, as it is a promising approach to overcome this drawback. It supports

¹ cakeflow.wi2.uni-trier.de

the search of a suitable query and hence the desired solutions by involving user interaction during adaptation.

Further, separating similarity-based retrieval and adaptation may provide workflows that can not be at best adapted according to the query. Hence, methods will be developed that also reflect the adaptability of the workflows during the retrieval stage [13].

Moreover, feedback of workflow adaptation will be captured in order to address the retaining of adaptation knowledge [5]. This is essential, as the quality of automatically learned adaptation knowledge can not always be ensured. Thus, the quality of workflow adaptation is improved. Further, the growth of adaptation knowledge can be controlled and hence the performance of adaptation can be maintained.

References

1. Aha, D.W., Muñoz-Avila, H.: Introduction: Interactive case-based reasoning. *Applied Intelligence* 14(1), 7–8 (2001)
2. Bergmann, R., Gil, Y.: Similarity assessment and efficient retrieval of semantic workflows. *Inf. Syst.* 40, 115–127 (Mar 2014)
3. Bergmann, R., Müller, G., Wittkowsky, D.: Workflow clustering using semantic similarity measures. In: Timm, Thimm (eds.) *KI 2013: Advances in Artificial Intelligence*. LNCS, vol. 8077, pp. 13–24. Springer (2013)
4. Hollingsworth, D.: Workflow management coalition glossary & terminology. http://www.wfmc.org/docs/TC-1011_term_glossary_v3.pdf (1999), last access on 04-04-2014
5. Jalali, V., Leake, D.: On retention of adaptation rules. In: *Case-Based Reasoning Research and Development*, pp. 200–214. Springer (2014)
6. Minor, M., Bergmann, R., Görg, S., Walter, K.: Towards case-based adaptation of workflows. In: *Case-Based Reasoning. Research and Development*, pp. 421–435. Springer (2010)
7. Minor, M., Montani, S., Recio-García, J.A.: Process-oriented case-based reasoning. *Information Systems* 40(0), 103 – 105 (2014)
8. Müller, G., Bergmann, R.: A Cluster-Based Approach to Improve Similarity-Based Retrieval for Process-Oriented Case-Based Reasoning. In: *Proceedings of ECAI 2014*. Prague, Czech Republic (2014)
9. Müller, G., Bergmann, R.: Workflow Streams: A Means for Compositional Adaptation in Process-Oriented CBR. In: *Proceedings of ICCBR 2014*. Cork, Ireland (2014)
10. Müller, G., Bergmann, R.: Generalization of Workflows in Process-Oriented Case-Based Reasoning. In: *28th International FLAIRS Conference*. AAAI, Hollywood (Florida), USA (2015)
11. Müller, G., Bergmann, R.: Learning and Applying Adaptation Operators in Process-Oriented Case-Based Reasoning. In: *Proceedings of ICCBR 2015*. Frankfurt, Germany (2015)
12. Müller, G., Bergmann, R.: POQL: A new query language for Process-Oriented Case-Based Reasoning (2015), to be submitted to LWA 2015 Trier, Germany
13. Smyth, B., Keane, M.: Retrieving adaptable cases. In: Wess, S., Althoff, K.D., Richter, M. (eds.) *Topics in Case-Based Reasoning*, *Lecture Notes in Computer Science*, vol. 837, pp. 209–220. Springer Berlin Heidelberg (1994)

Opinionated Explanations of Recommendations from Product Reviews

Khalil Muhammad

Insight Centre for Data Analytics,
University College Dublin, Belfield, Dublin 4, Ireland.
{firstname.lastname}@insight-centre.org

1 Introduction

Recommender systems are now mainstream and people are increasingly relying on them to make decisions in situations where there are too many options to choose from. Yet many recommender systems act like “black boxes”, providing little or no transparency into the rationale of their recommendation process [1]. Related research in the field of recommender systems has focused on developing and evaluating new algorithms that provide more accurate recommendations. However, the most accurate recommender systems may not necessarily be those that provide the most useful recommendations — due to the influence of how recommendations are presented and justified to users [2–4]. Therefore, recommender systems must be able to explain what they do and justify their actions in terms that are understandable to the user. An explanation, in this context, is any added information presented with recommendations to help users better understand why and how a recommendation is made [5]. Studies show that explanations help users make better decisions and are therefore provided for many reasons [6, 7], which normally align with the objective of the recommender system. Interestingly, explanations may sometimes be provided from the users (not from the recommender system) to justify their choices [8].

The availability of user-generated reviews that contain real experiences provides a new opportunity for recommender systems; yet, existing methods for explaining recommendations hardly take into account the implicit opinions that people express in such reviews even though studies show that users are increasingly relying on the reviews to make better choices [9]. Also, explanations usually provide a posthoc rationalisation for recommendations; but, this work is motivated by a more intimate connection between recommendations and explanations, which poses the question: *can the recommendation process itself be guided by structures generated to explain recommendations to users?*

This work builds on existing research in the areas of case-based reasoning, recommender systems and opinion mining to propose a novel approach for building explanations in recommender systems. We will also explore the potential of opinionated explanations in driving the recommendation process.

2 Research Plan

The core focus of this work is to explore the role of opinions in explaining recommendations. Accordingly, we have identified the following areas of interest:

Ranking, filtering and evaluating feature quality. Feature-level opinion mining algorithms that are capable of extracting very granular opinions, such as [10], yield noisy features because they rely on shallow natural language processing (NLP) techniques. Ultimately, these features lack context and are too fine-grained to be intuitive to users. For instance, it will be nonsensical to explain a hotel recommendation as *“because visitors liked the wire...”*, where ‘wire’ is a feature mined from reviews. Hence, the research question: *“how to rank, filter and evaluate features mined from reviews”*. We will use off-the-shelf opinion mining techniques but focus on developing methods for ranking features so that they can be filtered and evaluated for quality (i.e. the extent to which a feature is relevant and presentable to users in explanations). This involves creating new methods for summarising features so that only qualitative and comprehensible features are presented in explanations.

Generating opinionated explanations. Explanations normally demonstrate how one or more recommended items relate to a user’s preferences, normally through an intermediary entity such as a user, item, or feature. For instance, Netflix may use the movies that a user has rated highly in the past to explain a movie recommendation. And since user ratings are often unable to fully represent user preferences, there is a place of fine-grained opinions that are explicitly provided by users in textual reviews. We expect that explanations that are based on opinionated reviews will be more natural and convincing. Hence the research question: *how to use such opinions to generate explanations of product recommendations?*. We will use opinions from reviews to generate that justify a particular recommendation or sets of recommendations, and we will conduct live-user trials to test for its usefulness in decision-making.

Driving recommendations using explanations. To date, most recommender systems have treated explanations as an afterthought, presenting them alongside recommendations, but with little connection to the recommendation process itself. This work will explore the potential of using explanations to drive the recommendation process itself so that, for example, an item will be recommended because it can be explained in a compelling way. Hence the research question: *how to use explanations to support similarity metrics and ranking strategies in a recommendation process?*

3 Progress

To address the problem of feature quality, we used the approach in [10] to mine opinions from a dataset of TripAdvisor hotel reviews. Then, using various lexical and frequency-based filtering techniques, we removed noisy, less opinionated and unpopular features. The remaining features were summarised into higher-level representations by clustering them based on the words they co-occur with in

sentences of reviews. This feature representation allows us to replace a low-level feature (e.g. ‘orange juice’) with a more meaningful higher-level one (e.g. ‘breakfast’) that is suitable for use in explanations.

We developed a new method for generating personalized explanations which highlight the pros and cons of a recommended item to a user. Our approach focuses on the features that the user has mentioned in their reviews, and those mentioned about the recommended item by other users. In the explanation, we prioritize the features that are likely to be of interest to the user. Each feature is classified as a pro or con based on its sentiment, and it is ranked by its popularity with the user and the recommended item.

We also developed another explanation strategy that explains a recommended item in comparison with other recommendations. That is, the explanation presents features of the recommended item that are better or worse than its alternatives.

Acknowledgments. This work is supported by the Insight Centre for Data Analytics under grant number SFI/12/RC/2289.

References

1. Herlocker, J.L., Konstan, J.A., Riedl, J.: Explaining Collaborative Filtering Recommendations. In: Proceedings of the 2000 ACM Conference on Computer supported cooperative work, ACM (2000) 241–250
2. McNee, S.M., Riedl, J., Konstan, J.A.: Being Accurate is not Enough: How Accuracy Metrics Have Hurt Recommender Systems. In: CHI’06 extended abstracts on Human factors in computing systems, ACM (2006) 1097–1101
3. Pu, P., Chen, L., Hu, R.: Evaluating recommender systems from the users perspective: survey of the state of the art. *User Modelling and User-Adapted Interaction* **22**(4-5) (2012) 317–355
4. Knijnenburg, B.P., Schmidt-Thieme, L., Bollen, D.G.: Workshop on User-centric Evaluation of Recommender Systems and their Interfaces. In: Proceedings of the fourth ACM conference on Recommender systems, ACM (2010) 383–384
5. Tintarev, N., Masthoff, J.: Effective Explanations of Recommendations: User-Centered Design. In: Proceedings of the 2007 ACM conference on Recommender systems, ACM (2007) 153–156
6. Tintarev, N., Masthoff, J.: Designing and Evaluating Explanations for Recommender Systems. In: *Recommender Systems Handbook*. Springer (2011) 479–510
7. Zanker, M.: The Influence of Knowledgeable Explanations on Users’ Perception of a Recommender System. In: Proceedings of the sixth ACM conference on Recommender systems, ACM (2012) 269–272
8. Nunes, I., Miles, S., Luck, M., De Lucena, C.J.: Investigating Explanations to Justify Choice. In: *User Modeling, Adaptation, and Personalization*. Springer (2012) 212–224
9. Lee, J., Park, D.H., Han, I.: The Different Effects of Online Consumer Reviews on Consumers’ Purchase Intentions Depending on Trust in Online Shopping Malls: An Advertising Perspective. *Internet research* **21**(2) (2011) 187–206
10. Dong, R., Schaal, M., O’Mahony, M.P., McCarthy, K., Smyth, B.: Opinionated Product Recommendation. In: *Case-Based Reasoning Research and Development*. Volume 7969. (2013) 44–58

Integrated Maintenance with Case Factories for distributed Case-Based Reasoning Systems

Pascal Reuss^{1,2}

¹ German Research Center for Artificial Intelligence
Kaiserslautern, Germany
<http://www.dfki.de>

² Institute of Computer Science, Intelligent Information Systems Lab
University of Hildesheim, Germany
<http://www.uni-hildesheim.de>

1 Integrated maintenance with Case Factories

Maintenance approaches for Case-Based Reasoning (CBR) systems focus on a single CBR system and mainly on a single knowledge container like the case base, the adaptation knowledge or the similarity. There are a few approaches that deal with shifting knowledge between knowledge containers, especially between the case base and adaptation knowledge. These approaches have been successfully applied to CBR systems in the past. In many systems, especially multi-agent systems, the knowledge is distributed among several homogenous or heterogeneous knowledge sources. Therefore, a maintenance approach has to consider the dependencies between these knowledge sources as well as high-level maintenance goals. An example is removing one or more cases from a case base. Cases in other CBR systems could depend on one of the removed cases, so they may become inconsistent (to some degree). The system should suggest an appropriate maintenance action like removing the depending cases to keep the system's correctness/consistency. To address these dependencies between CBR systems and their knowledge containers, a new maintenance approach for distributed CBR systems is required. In the following, I will describe the idea of an integrated maintenance approach based on so-called Case Factories, that will be capable of generating a maintenance plan for multiple CBR systems, considering dependencies, and providing explanations for the generated maintenance actions. The approach is designed for multi-agent systems based on the SEASALT architecture ([2]), which supports distributed knowledge sources. To develop this approach, four research goals have to be reached:

- Extend the original Case Factory approach from single system support to multi system support.
- Integrate maintenance explanation capabilities into the new approach
- Develop a methodology to apply the new approach to existing multi-agent systems as well as integrating it into the development of new multi-agent systems
- Evaluate the new approach and the methodology within an industrial use case for a decision support system

1.1 Case Factory and Case Factory Organization

The original idea of the Case Factory (CF) is from Althoff, Hanft and Schaaf ([1]) and is based on the Experience Factory approach from software engineering. A CF consists of several software agents for different tasks like evaluating incoherence or modifying the case base. Each knowledge source, in this context CBR systems, has its own CF that is responsible for maintaining the dedicated knowledge. The original approach has to be extended to support the maintenance of the other three knowledge containers, namely vocabulary, similarity and adaptation knowledge, considering intra-system dependencies between the knowledge containers. The original approach contains several software agents to monitor the case-base and one agent to do the necessary maintenance actions. To support all knowledge containers reasonably some more agents are required to monitor these containers and the maintenance tasks should be split on several agents. An own maintenance agent per knowledge container is required to support parallel maintenance of the knowledge. Additionally, a supervising agent is required to coordinate the maintenance actions. This agent is also responsible for the communication between the multiple CFs. Monitoring a knowledge container means to notify the supervising agent about changes of the knowledge inside the container, e.g. removing a case, renaming a concept, or adding a rule. Evaluation of a knowledge container means to check a knowledge container for inconsistencies, problem solving competence or efficiency. Therefore, different evaluation strategies could be used, for example computing the coverage and reachability of cases ([9]) or inconsistency checking ([8], [7] and [6]).

To store information about dependencies, maintenance goals and evaluation criteria, a so-called Maintenance Map is used. The Maintenance Map is based on the Knowledge Map from Davenport and Prusak ([4]), which was adapted to multi-agent-systems by Bach et al. ([3]). In contrast to a Knowledge Map, the Maintenance Map is a bidirectional graph. The vertices represent the knowledge sources in a distributed knowledge-based system and the edges represent the dependencies between the single sources. The weights of the edges could be used to describe the importance of the dependency. Additional information like maintenance goals, metrics, thresholds, or constraints could be defined, too. The Maintenance Map can also contain information about the priority of maintenance actions or about associated maintenance actions. This may be helpful, when a combination of maintenance actions is necessary to preserve the competence or efficiency of a single CBR system or the MAS ([5]).

While a CF is able to maintain a single CBR system, a high-level Case Factory Organization (CFO) is required to coordinate the actions of all CFs and take the dependencies between the single CBR systems into account. This CFO consists of several additional software agents to supervise the communication between the CFs and the adherence of high level maintenance goals. Additionally, agents collect the maintenance suggestions from the CFs and derive a maintenance plan from all single maintenance suggestions. The agents are also responsible for checking constraints or solving conflicts between individual maintenance suggestions. In addition, a maintenance suggestion may trigger follow-up maintenance actions based on the dependencies between the CBR systems. The concept of the CFO allows to realize as many CFs and layers of CFOs as required. This way a hierarchy of CFOs can be established that is scalable and supports multi-agent systems with many agents and layers.

To support the knowledge engineer, the CF and CFO should explain, why a certain maintenance action was suggested. To give a CBR system explanation capabilities a lot of knowledge is necessary. The underlying research assumption here is that the minimal knowledge required for the explanation of maintenance actions is the same knowledge that is required to suggest a maintenance action. It follows, that the minimal knowledge for explanations already exists in the system, if the system is able to (reasonably) suggest maintenance actions.

1.2 Current state of research and research plan

Currently, the concepts of the CF and the CFO are defined. An initial version of a CF is implemented within a research system called docQuery in the travel medicine domain. This implementation has to be extended with a CFO to test the integrated maintenance approach. The next steps on conceptual level are to formalize the concept and advance the maintenance planning and the explanation strategies on a more detailed level and implement them, too. It is planned to complete the conceptual development of the integrated maintenance approach this year and implement the complete approach within the docQuery system. Next year, the approach will be applied to an industrial use case in the domain of aircraft diagnosis to test the approach within an industrial environment. The evaluation of the approach and methodology will be completed till the third quarter of 2016 and the dissertation will be written until the beginning of 2017.

References

1. Althoff, K.D., Hanft, A., Schaaf, M.: Case factory - maintaining experience to learn. *Advances in Case-Based Reasoning Lecture Notes in Computer Science* 4106/2006, 429–442 (2006)
2. Bach, K.: Knowledge Acquisition for Case-Based Reasoning Systems. Ph.D. thesis, University of Hildesheim (2013), dr. Hut Verlag Mnchen
3. Bach, K., Reichle, M., Reichle-Schmehl, A., Althoff, K.D.: Implementing a coordination agent for modularised case bases. In: *Proceedings of the 13th UK Workshop on Case-Based Reasoning*. pp. 1–12 (2008)
4. Davenport, T.H., Prusak, L.: *Working Knowledge: How Organizations Manage What they Know*. Havard Business School Press (2000)
5. Leake, D., Kinley, A., Wilson, D.: Learning to integrate multiple knowledge sources for case-based reasoning. In: *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*. pp. 246–251. Morgan Kaufmann (1997)
6. McSherry, D.: An adaptation heuristic for case-based estimation. In: *4th European Workshop, EWCBR-98*. pp. 184–195 (1998)
7. Racine, K., Yang, Q.: Maintaining unstructured case bases. In: *Second International Conference on Case-Based Reasoning*. pp. 553–564 (1997)
8. Smyth, B.: Case-base maintenance. In: *11th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems IEA-98-AIE*. pp. 507–516 (1998)
9. Smyth, B., Keane, M.: Remembering to forget: A competence-preserving case deletion policy for case-based reasoning systems. In: *Proceedings of the 13th International Joint Conference on Artificial Intelligence*. pp. 377–382 (1995)