

Semi-Distributed Development of Agent-Based Consultation Systems

Georg Buscher Joachim Baumeister Frank Puppe Dietmar Seipel
Institute of Computer Science, University of Würzburg, Germany
{buscher, baumeister, puppe, seipel}@informatik.uni-wuerzburg.de

ABSTRACT

This paper proposes a semi-distributed approach for the development of agent-based consultation systems. The key idea is: what can easily be done to enhance quality, and reduce redundancy is done centrally, while the mass of knowledge is acquired in a distributed way.

Categories and Subject Descriptors

I.2.1 Applications and Expert Systems, I.2.11 Distributed Artificial Intelligence, H.3.3 Information Search and Retrieval

Keywords

knowledge acquisition, collaborative development, agent-based consultation systems, large-scale knowledge-based systems

1. INTRODUCTION

An agent-based consultation system consists of a society of agents offering consulting services to the user. It emphasizes the active role of the user during the problem solving process, since the user can select which agents to consult. While each agent delivers a useful consultation service on its own, they are designed to complement each other, so that they together may offer consultation services comparable to traditional monolithic consultation systems (MCS; see e.g. [1]). An example for a user-centered consultation system (UCCS) for diagnosis is presented in [2]. It allows for distributed knowledge acquisition from volunteer contributors [3].

Both UCCS and MCS have advantages and problems concerning knowledge acquisition. MCS suffer from the bottleneck of centralized knowledge acquisition making it nearly impossible to develop large systems. However, in particular in medical domains, where symptoms are usually highly ambiguous, large systems can provide a much higher benefit to the user than specialized systems. Large MCS suffer from the rigidity of the chosen knowledge representation, which usually does not fit all parts of the system equally well. A uniform knowledge representation also causes inflexibility concerning user interaction; adapting the knowledge for different user types further increases the already nearly prohibitive knowledge acquisition effort. UCCS on the other hand allow for a distributed development, since each agent

has well defined interfaces to its environment (i.e. the other agents). There is a structured interaction in UCCS among the agents: For example, a symptom class agent suspects diagnoses based on entered symptoms, clarification agents validate a diagnosis suspected by symptom class agents or by the user and a therapeutic agent determines adequate therapies for a validated diagnosis. The agents may have knowledge representation suitable for their purposes. Another advantage of UCCS compared to MCS is that the system is already usable with only some agents available. However, the knowledge of the different agent types overlaps, because a single symptom-diagnosis relation is usually needed by a symptom class agent as well as by a clarification agent resulting in redundancy problems. In addition, the distributed development makes it difficult to ensure quality standards over all agents. Further on, some services of a MCS are difficult to achieve by a UCCS like systematic dialogue guidance enabling high quality documentation and subsequent data mining options.

As a result a combination of centralized and distributed development may help to overcome some of the weaknesses of the respective approaches. The key idea is: what can easily be done to enhance quality, and reduce redundancy is done centrally, while the mass of knowledge is acquired in a distributed way. The “knowledge champion” initializes the consultation system by providing a basic list of necessary agents and may suggest or define the coarse structure of the agents. He or she is also responsible for organizing quality control measures (e.g. based on peer review or feedback from users consulting the UCCS), initiating activities to overcome detected weaknesses and decide in controversies resulting from the distributed process.




Figure 1. Development tasks and their effort.

Most of the concepts described in the following are exemplified by a pub recommendation system (PRS). It contains both a hierarchy of selection criteria (symptom classes) and of pubs (diagnoses). To every selection criterion, a formal symptom class agent is assigned, to every pub, a formal and an informal clarification

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Conference'04, Month 1–2, 2004, City, State, Country.
Copyright 2004 ACM 1-58113-000-0/00/0004...\$5.00.

agent exist. The development tasks accomplished in centralized and distributed development are presented in Figure 1.

2. ASPECTS OF SEMI-DISTRIBUTED DEVELOPMENT

There are several important issues concerning semi-distributed development of a UCCS being described in the following:

The *knowledge representation of the UCCS agents* has to be suitable for distributed development. In general, set-covering models and Bayesian Nets seem to be appropriate here due to the independency of the knowledge concerning different diagnoses. Heuristic rules on the other hand are more problematic for distributed development because the collaboration of the different rules concerning several diagnoses is a very sensitive process. For the formal agents of the PRS, for example, we chose set-covering knowledge, because the specification of set-covering relations between pubs and pub characteristics is a relatively simple process.

Avoiding redundancy of symptom class and diagnosis clarification agents: As stated in the introduction, the knowledge of the different agent types overlaps, because the same symptom-diagnosis relation is usually needed by a symptom class agent as well as by a clarification agent. The general idea to avoid redundancy is to generate the knowledge of symptom class agents from the knowledge of diagnosis clarification agents. This requires using the same terminology for both agent types. We achieve this in the PRS by predefining a centralized terminology (see Fig. 1; a task of the knowledge champion). Thus different developers add knowledge only for diagnosis clarification agents in a distributed but coordinated way.

Omitting *terminology alignments between different agents* is a further advantage of using a centralized terminology for all agents. Otherwise, due to the requirement that the agents should be able to exchange meaningful data between each other, alignments between the different terminologies of different agents would be necessary and this normally includes problems concerning a loss of precision.

Modifications of the centralized terminology might affect all agents in the system. Removing a symptom involves removing the knowledge of all agents concerning this symptom. Adding a new symptom requires the insertion of default knowledge for this symptom to all agents. With respect to the PRS, this implies inserting zero set-covering relations concerning the different pubs and the new symptom, which means that the new symptom neither counts for nor against a pub. In addition, an alerting system is provided, which informs the agents' developers about such modifications.

Knowledge acquisition with the help of flexible templates can be applied to reduce inconsistencies of the distributedly entered knowledge and to simplify the knowledge acquisition process [4]. Concerning the PRS, for example, the developers can put in the characteristics of a new pub with the help of web-templates automatically generated from the centralized terminology. Thus, more people can be motivated to contribute to the system. Furthermore, wrong or inconsistent inputs can be avoided, which benefits the quality of the whole system.

Multiple opinions of the agents' developers are a typical problem in distributedly developed systems. There are mainly three alterna-

tives, how to handle this problem: First, the strict way would be to allow only one agent for every entry (i.e. symptom class or diagnosis), which cannot be modified by other developers than its author. Developers of other opinions only have the possibility to add comments for the agent. Second, in addition to the first alternative, every developer could be allowed to modify an agent. This might lead to faster development but might also yield unproductive controversies as they can sometimes be observed in www.wikipedia.org. Third, in addition to the first and second alternative, the developers could be allowed to add a new agent for the same entry in parallel stating a different opinion, which means that inconsistencies between agents for the same entry are explicitly tolerated. Currently, we experiment with all three alternatives in the PRS.

Feedback management is important to improve the quality of the different agents and therefore the whole system. The users have the possibility to give explicit feedback by rating an agent with a grade or writing a comment. This feedback is visible to all users and in particular to the authors of the respective agents, who might change the agents accordingly. In this way, an iterative development cycle emerges.

3. DISCUSSION

The proposed work can be viewed as a special realization of the semantic web idea. Instead of developing HTML- or Wiki-pages for pub recommendation directly, such pages are generated with a specific underlying semantic via knowledge templates and predefined terminologies in a distributed process with volunteer contribution. This knowledge is used by an inference engine and for generating a user-friendly search interface to increase precision and recall in comparison to standard search engines. We plan, if parts of the knowledge are already available in the Web (e.g. prices for beverages in the PRS), the knowledge base should be updated automatically by information extraction agents. Since the areas of application for UCCS reach from medical and legal advice via product selection to most diagnostic and classification domains, we also started other applications (e.g. MediSuggest for medical advice; www.medicococonsult.de; in German).

4. REFERENCES

- [1] Huettig, M., Buscher, G., Menzel, M., Scheppach, W., Puppe, F., Buscher, H.-P.: A Diagnostic Expert System for Structured Reports, Quality Assessment, and Training of Residents in Sonography: Medizinische Klinik 99, 117-122, 2004
- [2] Buscher, G., Baumeister, J., Puppe, F., Seipel, D.: User-Centered Consultation by a Society of Agents. Proc. 3rd International Conference on Knowledge Capture (K-CAP 2005), Banff, Canada, 27-34, 2005
- [3] Symposium on Knowledge Collection from Volunteer Contributors (KCVC-05), AAAI Spring Symposium, technical report SS-05-03, <http://www.aaai.org/Library/Symposia/Spring/ss05-03.php>, 2005
- [4] Chklovski, T., Gil, Y.: Improving the Design of Intelligent Acquisition Interfaces for Collecting World Knowledge from Web Contributors. Proc. K-CAP 2005, Banff, 35-42, 2005