TOWARDS MULTIAGENT INTERACTIVE ADAPTION OF INDIVIDUAL MODELS FOR DECISION SUPPORT*

Tobias Warden, Thomas Wagner, Hagen Langer and Otthein Herzog
Center for Computing and Communication Technologies (tzi), Universität Bremen, Bremen, Germany

Keywords: Multiagent systems, Interaction protocols, Distributed knowledge management, Machine learning, Knowledge exchange, Interactive learning.

Abstract: Software agents in complex, dynamic environments need to update, adapt, and improve their knowledge models for decision making in order to achieve adequate results. Their individual adaption often relies on machine learning from observational data. However, when data is not available in the required quantity and quality, alternative approaches are required. We propose an interaction-based approach to individual model adaption in multiagent systems, describe agent roles and interaction principles, and discuss how a goal-oriented transfer of knowledge among agents can be integrated into an agent-based knowledge management framework.

1 INTRODUCTION

Besides adequate access to decision-relevant information and domain knowledge, access to tailored dynamic knowledge is a competitive advantage for the performance of intelligent agents in complex task domains such as health care, supply chain management, electronic markets, or autonomous logistics. Referring to *dynamic knowledge*, we concur with Jennex who defines *knowledge* as information, ”understood such that it explains the how and the why about something” (Jennex, 2009), i.e., patterns applicable in situation assessment and prediction. Dynamic knowledge subsumes several categories of integral knowledge models for decision making, including decision models, prediction models, or classification models.

Depending on the specific use case, the respective models may be placed at an agent's disposal on instantiation. This means that the models have been either learned offline from historic data or handcrafted in a knowledge engineering process. In other cases, the agents themselves may be equipped with the prerequisites for online learning of models, based on individual experience or accessible data repositories. We consider agents which not only compile and employ classification models in their primary domain roles but also possess the ability to assess the performance of their models. This kind of meta-reasoning in particular enables the agents to monitor their performance over time and determine when an adaption of momentarily operationalized classification models is necessary.

However, left to its own devices, it is often not possible for a deployed agent to effectuate the desired model adoption. For instance, the necessity for adaption may allude to a preset model, provided for use by an external source such as a knowledge engineer. Consequently, the agent may not have access to a representative pool of data to address adaption by means of individual re-learning on this data. In other cases, the ability to relearn may exist but the costs of a new learning sweep are prohibitive or it has become evident that available learning schemes cannot provide an improved model. Thus, while self-assessment of an agent reveals a necessity to adapt a decision support model, its skill set and/or resources inhibit a self-sufficient adaption.

We claim that the multiagent environment provides the key to complement the individual adaption capability of an agent. Within a persistent multiagent system (MAS), there are often multiple agents with kindred primary tasks and associated decision support models. *These models* may contain knowledge that is both relevant and useful for an agent with an adaption deficiency. We propose a knowledge management (KM) framework which allows agents to tap that knowledge in a goal-oriented way. Our approach

*This research has been supported by the German Research Foundation (DFG) within the Collaborative Research Centre 637 ”Autonomous Cooperating Logistic Processes – A Paradigm Shift and its Limitations” at the Universität Bremen, Germany.

Warden T., Wagner T., Langer H. and Herzog O..
TOWARDS MULTIAGENT INTERACTIVE ADAPTION OF INDIVIDUAL MODELS FOR DECISION SUPPORT.
DOI: 10.5220/0003886905530559
In *Proceedings of the 4th International Conference on Agents and Artificial Intelligence* (IWSI-2012), pages 553-559
ISBN: 978-989-8425-95-9

builds upon dedicated knowledge management functions executed by both the agent seeking model adaption and its supporters. It enables individual adaption in knowledge networks. Specifically, it acknowledges that dynamic knowledge may not be simply transferred among agents. Rather, interactive adaption as we see it requires that the agents seeking adaption actively interpret and integrate advice from peers, thus effectively re-constructing individual advisor knowledge. Thus, our approach to address adaption deficiencies through social interaction emulates, in spirit rather than cognitive accuracy, human problem solving strategies. Also, instead of imposing a fixed organizational structure, the approach is fully distributed. The framework builds upon and extends approaches pursued in the machine learning (ML) community, such as active learning or transfer learning.

The following section first provides a detailed problem analysis and relates our methodology to approaches from the ML community. Section 3 then adds a knowledge management and multiagent perspective. It introduces our framework that enables interactive adaption of individual decision support models. Section 4 highlights related research, then conclude with a discussion and future work.

2 PROBLEM ANALYSIS

The minimal common prerequisites for agents that we seek to support through interactive model adaption are: 1) For competitive action in their primary domain role, the agents rely on a classification model. We make no assumptions as to the origin of this model, although models learned by the agent itself are our primary concern. 2) The agents actively monitor both performance of their models and any self-sufficient model adaption. Based on these prerequisites, a range of scenarios for our approach can be spanned.

The first distinction refers the embedding of a learning subsystem. Agents *without* this feature cannot adapt their model through learning from training data. For adaption, they are thus *necessarily* dependent on external help. For agents *with* a learning subsystem, a second distinction refers to whether individual learning constitutes a feasible option alternative in the momentary context. Engaging in learning may be infeasible due to the following reasons: 1) Training data may not be accessible, at all, or be insufficient. 2) The costs in terms of either computation/memory resources or computation time incurred by (re-)learning a model may be prohibitive. In both cases, self-sufficient learning is not *practical* as means to improve a model. Hence, agents in these situations have an interest in enlisting external help for a target-oriented direct model modification. Even if (re-)learning was a feasible option in principle, agents may still have to face problems: The available data constitutes a skewed sampling of an underlying true distribution and is thus not representative or data may only allow to learn a concept description which does not comply with desired quality characteristics. For classification, this would be objectives with respect to acceptable precision and recall, or false-positive/false-negative classifications. While learning is possible in these cases, due to the input data it does not provide a model which complies with preset performance standards or which bests the currently operationalized model. Therefore, agents in such situations may have an interest in enlisting external help designed to either broaden their data basis or acquire advice which essentially allows their machine learning scheme to derive additional value from an unvaried data basis.

We have shown scenarios where agents can draw a surplus from a flexible framework for interactive adaption of their models. An additional important finding is that the required complementary adaption skill sets vary depending on the respective situation. From a machine learner's perspective, selected problem instances laid out in the problem analysis have been addressed in the literature. Two strands of ML-research with special relevance are *transfer learning* from previous learning tasks an *interactive learning* supported by other individuals.

Transfer learning addresses the specific problem of an insufficient data basis to learn a model through re-use of training data originally collected for related learning tasks (for a survey on transfer learning for classification problems, see (Pan and Yang, 2009)). A broadening of the data basis is also a primary goal in the area of *active learning* (Settles, 1995). The approach has been suggested amongst others for classification problems where the labeling of training instances for learning is costly or time-consuming. Beginning with a small pool of labeled training data, the learning system iteratively learns a model and assesses which additional data would provide the best chance to optimize the hitherto learned model. This data is then made available through intervention of a human domain expert or acquired autonomously, e.g., through conduct of experiments.

Možina et al. have proposed *argumentation based machine learning* (ABML). This approach allows to attain improvements in the performance of a learning system which needs to solve a classification task by a human domain expert as an interaction counterpart (Možina et al., 2007). The learning system is endowed with self-assessment capabilities in that it monitors its own learning progress and, specifically, identifies problem instances in the training data, that are particularly ill-covered by its learned concept description. The learner then reaches out to its preset human interaction partner, presenting these instances as queries. The expert then uses his domain knowledge to provide a machine-readable explanation (called *argumentation*). These are accounted for in subsequent learning phases.

ABML addresses what we consider the most challenging flavor of interactive model adaption in that it tries to tap on the implicit domain expertise of an advisor to augment the existing advisee training data, thus enabling learning progress. However, the approach needs to be elevated to multiagent learning. Here, other learning agents as artificial domain experts would then subsume the single human expert involved in ABML. Since specific strengths of ABML that have been highlighted by Možina et al. include not only "*effective reconstruction of expert's prior knowledge*" (Možina et al., 2007, p. 932) and favorable results even in situations where the human interaction partner's "*knowledge about the domain is superficial (might be wrong occasionally)*" (Možina et al., 2007, p. 933), this approach seems promising.

3 A MULTIAGENT APPROACH FOR INTERACTIVE ADAPTION

The methodical assessment of robust interactive adaption and, in the process, transfer of individually learned knowledge among cooperating agents in multiagent systems requires a suitable context from a knowledge management point of view. We draw on a framework for distributed knowledge management that we have originally proposed for intelligent agents jointly realizing control of autonomous logistic processes (Langer et al., 2006). We will use and extend this framework as a point of origin to derive necessary knowledge management roles, means for interoperability, and intra- as well as inter-agent organization of multiagent adaption.

3.1 Role-based Distributed Knowledge Management in MAS

This framework focusses on KM whose functions are *procured by* software agents. More importantly, these functions are also *designed for use by* agents as artificial decision-makers. In complex task environments, e.g., autonomous logistics, the availability of diverse knowledge has been identified as a key factor for an effective treatment of the primary agent roles in the modeled processes with a desired quality of service (Gehrke et al., 2010). Often, initial provision of such default knowledge alone is not sufficient in order to accommodate for the complexity and dynamics of the task environment. It then becomes necessary to design adaptive agents, capable of individual knowledge revision and the compilation of tailored models via learning. Over time, knowledge hence becomes to an increasing degree tailored to its task context. Thus, analogously to the situation with employees within organizations, knowledge is spread rather than accumulated in a centralized knowledge repository as assumed in conventional knowledge management approaches.

This situation is specifically accounted for in the KM framework through encapsulation of well-differentiated knowledge management functions as agent *roles*. Examples proposed in (Langer et al., 2006) cover amongst others *knowledge acquisition*, *knowledge processing* or *brokerage* (See Figure 1). The strength and flexibility of the role abstraction is that, in contrast to other approaches (Van Elst et al., 2004), knowledge management-related abilities are not restricted to highly-specialized dedicated agents. Rather, any agent is free to assume a time-variant set of knowledge management roles as deemed appropriate in its situation context. These roles are understood as *auxiliary roles* which complement *domain-specific primary roles* such as the management of logistic processes. The roles can be further categorized into *internal* and *external roles*. We adopt a notion of internal roles where these are characterized 1) by *reasoning capabilities* and 2) a *deliberation pattern* (i.e., a plan to accomplish the knowledge management task). Internal roles can be conducted self-sufficiently. External roles by contrast require interaction, structured by one or more *interaction protocols*.

3.2 Knowledge Management Roles for Interactive Adaption

Two complementary roles are involved in interactive adaption, namely the *advisee role* played by an agent seeking assistance in adaption and the *advisor role* played by a temporary domain expert. In our multiagent environment, these roles are no longer tied to specific types of actors. We now introduce additional knowledge management roles and role aspects (See Figure 1) which extend our framework for distributed knowledge management (Langer et al., 2006).

The *model acquisition role* is a specialization of the *knowledge processing role*, it presupposes access

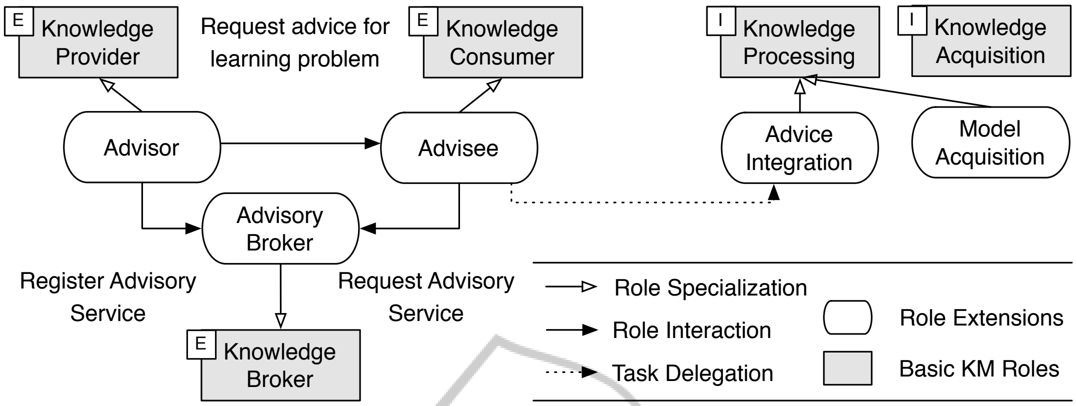

Figure 1: Adopted roles in the distributed KM framework by Langer et al. (Langer et al., 2006) (gray) and extensions for interactive adaption of individual models.

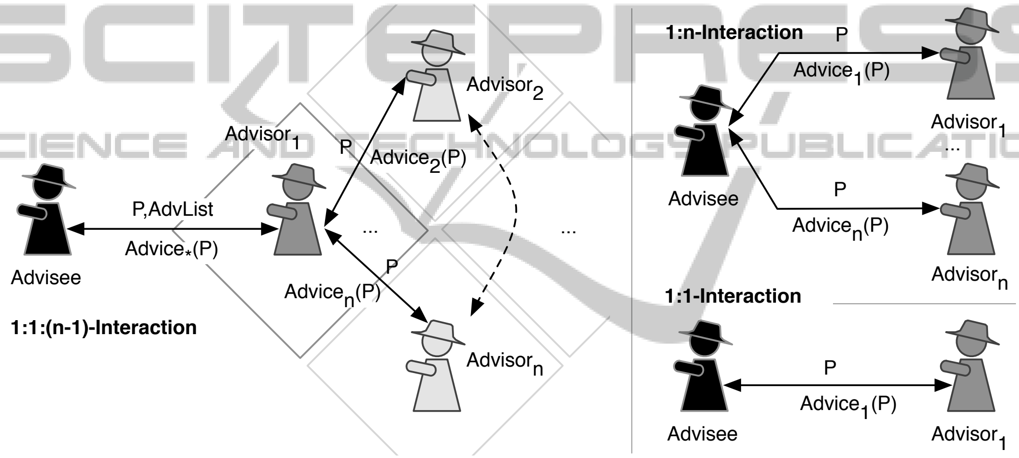

Figure 2: Overview of possible interaction patterns for knowledge transfer episodes.

to representative training data and an adequate machine learning scheme (e.g., a decision tree learner). Contingent on the agents' primary domain role(s), the data used for machine learning may constitute individual experience gathered through action in the task domain. It may as well originate from a data repository accessible to the learner. Once a model such as a classifier has been learned successfully, we assume that the role also exposes its inferential capabilities for internal used by the agent.

The *advisee role* is a specialization of the *knowledge consumer role*, elevated to the level of knowledge rather than information. Any agent may assume the advisee role when an assessment of its decision support model has shown deficiencies in the model performance that cannot be handled by internal means alone. Such a situation may arise due to an insufficient access to representative training data such that the prerequisites for the *model acquisition role* cannot be met. If the agent is keen to improve its model immediately rather than wait for new data to become accessible, the advisee role offers a viable action alternative. The agent then becomes an active learner in that it actively seeks for and eventually approaches peers that assume the learning advisory role introduced below. In the interaction associated with these roles, advisors are presented with learning problems and asked to offer advice to address said problem based on their models.

The *advice integration role* is a specialization of the knowledge processing role. Advice integration is understood as a subsidiary task to succeed with the *advisee role*. We propose that the superordinate advisee role delegates advice integration to this internal role. One can conceive different feasible interpretations of this role. As a first option, the advice provided as input may be used to directly revise an existing model (e.g., by pruning or expanding branches in a decision tree, or revision of a rule set). A second option is to conceive the advice integration as a specialization of the *model acquisition role* presented above. In such a case, a new model is learned based on the initial training data and the accumulated advice as additional background knowledge to bias/focus the operation of a learning scheme that is able to handle the additional input.

The *advisor role* is the role complement to the *advisee role* and a specialization of the *knowledge provider role*. It may be played by any agent with access to a particular decision support model if it wants to provide an knowledge advisory service for agents which need to learn and maintain similar models. The advisor model constitutes a valuable asset in that it captures relevant domain expertise for agents in similar application contexts. The model which is used as basis for the advisory service may be hand-crafted, yet the probably more interesting use case involves individually learned models. In order to play an advisor role, it is necessary to interpret requests placed by advisees and compile tailored advice to address the communicated learning subproblem, based on the advisor's own model. The role abstracts from the particular type of decision support model used by an advisor. While the adoption of an advisee role is an immediate consequence of a concrete need, the complementary advisor role may be played persistently.

The *advisory broker role* is an external role. It is devised as a specialization of the *knowledge broker role*. An agent in this role acts as a specialized yellow pages service within the MAS. It administers meta-information about knowledge advisory services exposed by agents currently assuming one or more advisor roles for specific learning problems. The meta-information is deposited by the advisors. It specifies amongst others the respectively handled learning task (e.g., classification), any bias towards particular learning objectives (e.g., avoidance of false positive classifications) and a meta-description of the learning domain. An advisory broker also accepts requests by agents seeking advice for a learning problem. The broker matches the request information against its advisory portfolio in order to point the requesting agents to suitable interaction partners.

3.3 Structuring Agent Interaction for Knowledge Transfer Episodes

Further building blocks for our approach are means to organize interactive adaption. To that end, we propose a two-layer approach: On the lower layer, an interaction protocol is used to structure the course of a single *knowledge transfer episode*. Each such episode spans the acquisition of a single piece of advice relating to a specific learning subproblem. On the upper level, agents need to structure their global interactive learning process which may involve multiple consecutive lower-level transfer episodes. Each is thereby characterized by a specific learning subproblem which is brought up by an advisee. In the episode, the advisee engages in direct interaction with one or more suitable advisors which need to be known beforehand. Given such a list of advisors, the advisee can then utilize several interaction patterns as sketched in Figure 2.

Single-Tier 1:1 Interaction: As a first option, the advisee may choose one advisor for exclusive interaction. As a result, the advisee will receive a single piece of advice to resolve the learning subproblem addressed in the transfer episode. If the advisor refuses to advise on the specified problem, the outcome may also be a failure of the learning episode.

Single-Tier 1:n Interaction: In order to broaden advice acquisition for the learning subproblem at hand, and at the same time reduce to the failure potential of the transfer episode, the agent may place the same advisory request with a larger number of advisors. As a result of this interaction pattern the advisee may be able to receive multiple independently compiled pieces of advice. As a consequence, it lies within the responsibility of the advisee to perform advice integration. In case of complementary advice, such an operation may be straightforward. However, advice integration also involves conflict resolution for contradictions. An interaction protocol for this interaction pattern is depicted in Figure 3.

Multi-Tier 1:1:n Interaction: The two interaction patterns discussed so far are single-tier interactions in which the advisee interacts directly with all advisors involved in the transfer episode. An alternative is a multi-tier interaction which again involves a group of advisors. In contrast to the preceding interaction patterns, the advisors appear as a *holon*[2] with a single advisor acting as *holon head*. The advisory holon is created dynamically upon request by the advisee. Besides the learning subproblem that constitutes the topic of the interaction, the advisee also communicates to an initial advisor, acting as holon head, a set of additional advisors. These then are to be consulted while preparing a single, consolidated piece of advice. The initial advisor uses this information about the additional advisors to relay the learning problem. It is also responsible for advice integration. The interaction among the advisors can itself follow different interaction patterns depending on the preferred method

[2] In the categorization of holonic MAS by Fischer et al. (Fischer et al., 2003), the organizational form of an *moderated association* is adequate.

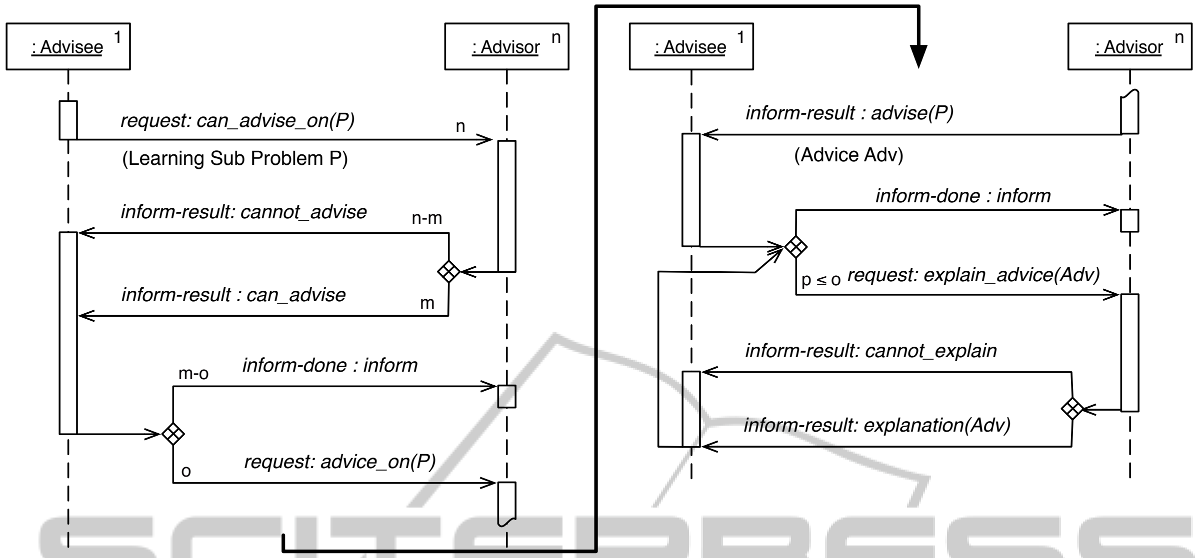

Figure 3: Protocol for single-tier 1:n interaction in a knowledge transfer episode.

of advice integration.

A Single-Tier 1:n Interaction Protocol We now present an interaction protocol for the single-tier 1:n interaction pattern (See Figure 3). The protocol consists of multiple stages. First, the advisor requests from all potential advisors, whether or not they consider themselves fit to provide advice for the learning subproblem at hand. The advisee waits for the respective advisor responses. Based on this feedback, the advisee can select the subset of advisors from which it can then request actual advice. In case the advisors do not only communicate a simple binary decision with respect to their willingness to advise, but also include information about their confidence to provide '*good*' advice, the advisee may choose to continue the conversation with a selection of all capable advisors. Those which are excluded are consequently informed that their services are no longer required. The remaining advisors receive requests to advise. The interaction protocol in Figure 3 then envisages that the earlier consent to advise enforces that each advisor actually provides individual advice and does not back out at this stage of the conversation. In case the communicated advice proves to be comprehensible for the advisee, the interaction with the respective advisor ends here. If, however, the advisee cannot comprehend any piece of advice directly, the interaction protocol provides for a continuation of the conversation to (repeatedly) request additional information to further understanding.

4 RELATED WORK

Interactive adaption of decision support models bears a cross-sectional character with related work from several research fields, namely machine learning (ML) and multiagent systems (MAS). In Section 2, we introduced related ML-research, most notably *argument based machine learning* by Možina et al. (Možina et al., 2007). While still restricted to adaption of rule-based classification models (AB-CN2), the approach has been embraced by Napierala and Stefanowski for the MODLEM rule induction algorithm (Napierala and Stefanowski, 2010). We consider these approaches a vantage point to implement the *advice integration role*. Ontañón and Plaza propose an approach to interactively resolve inconsistencies in individually learned concept descriptions based on computational argumentation (Ontañón and Plaza, 2010). As their goal is the consolidation of individual models, the approach is interesting when advice from multiple advisors must be consolidated to reach consensus. Costantini and Tocchio propose an approach for learning by knowledge exchange in logical agents (Costantini and Tocchio, 2005). The authors concentrate on the exchange of agent beliefs and rules encoding action recipes. They contemplate the role of trust in exchanging knowledge and propose strategies for graded operationalization of acquired knowledge. Jakob et al. implement adaptability in MAS by means of collaborative logic-based learning, focussing on communication strategies for acquired knowledge (Jakob et al., 2008). At the cross-section of MAS, robotics, and reinforcement learning, the implementation of social learning strategies have drawn significant attention (see, for instance, (Noble and Franks, 2003)).

5 DISCUSSION AND FUTURE WORK

In this paper we have sketched an interaction-based approach to the adaption of individual decision support models in MAS. It is desirable when software agents in complex, dynamic environments need to update, adapt, or improve their knowledge base for decision making. Sometimes, this improvement process can be based on machine learning from observational data, alone. But when available data is insufficient in quantity or quality, when data is too expensive, or when the machine learning process turns out to be too complex, alternative approaches are needed. There are two basic components in our approach: 1) a set of specific knowledge transfer roles which extends a set of basic knowledge management roles, and 2) a collection of interaction protocols for knowledge transfer.

Ongoing work on the proposed multiagent framework comprises a prototype implementation of the KM roles and interaction patterns introduced in Section 3. We focus first on the adaption of rule-based classification models, relying, for advice integration, on the ABML approach by Možina et al. In the process, we will also elaborate the meta-control, used by an advisee to guide its interactive adaption process over multiple knowledge transfer episodes, as a flavor of local search in a model space. The prototype is implemented based on the *JADE* agent development environment. Evaluation will be performed in the *PlaSMA* multiagent-based simulation environment (Warden et al., 2010). In the future, we also seek to enable a more far-reaching interoperability between heterogeneous agents in the context of interactive model adaption. This includes support for diversity in employed models (e.g., rule-based for the advisee and ANN-based for the advisor(s)). It also includes support for heterogeneity in the training data (with respect to attributes), discretization of values in individual learning, and the naming of attributes and concept classes. These extensions specifically call for the provision of additional KM roles, enabling, for instance, semantic mediation.

REFERENCES

Costantini, S. and Tocchio, A. (2005). Learning by Knowledge Exchange in Logical Agents. In *WOA*, pages 1–8. Pitagora Editrice Bologna.

Fischer, K., Schillo, M., and Siekmann, J. (2003). Holonic Multiagent Systems: A Foundation for the Organisation of Multiagent Systems. In *Holonic and Multi-Agent Systems for Manufacturing*, pages 1083–1084. Springer.

Gehrke, J., Herzog, O., Langer, H., Malaka, R., Porzel, R., and Warden, T. (2010). An Agent-Based Approach to Autonomous Logistic Processes. *Künstliche Intelligenz*, 24(2):137–141.

Jakob, M., Tožička, J., and Pěchouček, M. (2008). Collaborative Learning with Logic-Based Models. In *Adaptive Agents and Multi-Agent Systems III. Adaptation and Multi-Agent Learning*, volume 4865 of *LNCS*, pages 102–116. Springer.

Jennex, M. (2009). Re-Visiting the Knowledge Pyramid. In *Proceedings of the 42nd Hawaii Intl. Conference on System Sciences, 2009.*, pages 1–7. IEEE.

Langer, H., Gehrke, J. D., Hammer, J., Lorenz, M., Timm, I., and Herzog, O. (2006). A Framework for Distributed Knowledge Management in Autonomous Logistic Processes. *Intl. Journal of Knowledge-Based & Intelligent Engineering Systems*, 10(4):277–290.

Možina, M., Zabkar, J., and Bratko, I. (2007). Argument Based Machine Learning. *Artificial Intelligence*, 171(10-15):922–937.

Napierala, K. and Stefanowski, J. (2010). Argument Based Generalization of MODLEM Rule Induction Algorithm. In *Rough Sets and Current Trends in Computing*, volume 6086 of *LNCS*, pages 138–147. Springer.

Noble, J. and Franks, D. W. (2003). Social Learning in a Multi-Agent System. *Computing and Informatics*, 22(6):1001–1015.

Ontañón, S. and Plaza, E. (2010). Concept Convergence in Empirical Domains. In *Proceedings of the 13th Intl. Conference on Discovery Science*, pages 281 – 295.

Pan, S. and Yang, Q. (2009). A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, pages 1345–1359.

Settles, B. (1995). Active Learning Literature Survey. *Science*, 10(3):237–304.

Van Elst, L., Dignum, V., and Abecker, A., editors (2004). *Agent-Mediated Knowledge Management: International Symposium (AMKM '03)*, Standford, CA, USA. Springer. Revised and Invited Papers.

Warden, T., Porzel, R., Gehrke, J., Herzog, O., Langer, H., and Malaka, R. (2010). Towards Ontology-based Multiagent Simulations: The PlaSMA Approach. In *Proc. of the 24th European Conference on Modelling and Simulation*, pages 50–56.