

Situation-aware Risk Management in Autonomous Agents*

Martin Lorenz, Jan D. Gehrke,
Hagen Langer, and Ingo J. Timm
Center for Computing Technologies – TZI
University of Bremen
Am Fallturm 1, 28359 Bremen
(mlo|jgehrke|hlang|j.timm)@tzi.de

Joachim Hammer
Dept. of Computer & Information Sciences &
Engineering
University of Florida
Gainesville, FL 32611-6125
jhammer@cise.ufl.edu

ABSTRACT

We present a novel approach to enable decision-making in a highly distributed multiagent environment where individual agents need to act in an autonomous fashion. Our architecture framework integrates risk management, knowledge management, and agent deliberation to enable sophisticated, autonomous decision-making. Instead of a centralized knowledge repository, our approach supports a highly distributed knowledge base in which each agent manages a fraction of the knowledge needed by the entire system.

Categories and Subject Descriptors: I.2.11 [Distributed Artificial Intelligence]: Intelligent agents, H.1.1 [Systems and Information Theory]: Value of information, H.4.2 [Types of Systems]: Logistics, Decision support

General Terms: Design, Economics, Management, Reliability

Keywords: Risk management, knowledge management, decision support, intelligent agents, logistics

1. INTRODUCTION

In this paper we describe an approach to enable robust decision-making in a highly distributed, multiagent environment where agents need to act in an autonomous fashion. Our application is the logistics domain where autonomous agents are seen as a promising and effective approach to represent the different planning, scheduling, and controlling processes in an enterprise. For example, we can envision a scenario in which agents are used to represent real-world entities such as trucks and containers, abstract objects such as weather or traffic services, or even human decision makers, such as a ramp agent at a loading dock.

In this scenario, the agents need to make decisions about which containers to transport, what the fastest route to a specific destination is given current road or weather conditions, or what to do with goods damaged during unloading, for example.

Enabling this type of autonomous decision-making is challenging given the potentially large number of agents that could be involved, and the dynamic and sometimes even competitive environment in which the agents operate. In the context of this work we define risk as uncertainty about the future state of the world which implies that any decision by the agent might turn out wrong. The goal of risk management (RM) is to attempt to optimize the agent's

decisions in the presence of incomplete, imprecise, or debatable information by reducing the uncertainty about future events.

Knowledge management (KM) is an important means to achieve this. Our approach to KM aims to find a rational basis to obtain the needed information. Furthermore, the agent is challenged by the fact that the knowledge that is needed is often highly dynamic, context-sensitive, incomplete, or uncertain. Thus, the integration of risk and knowledge management enables context-based, situation-aware decision-making, which in turn supports autonomous, self-managing behavior of the agents.

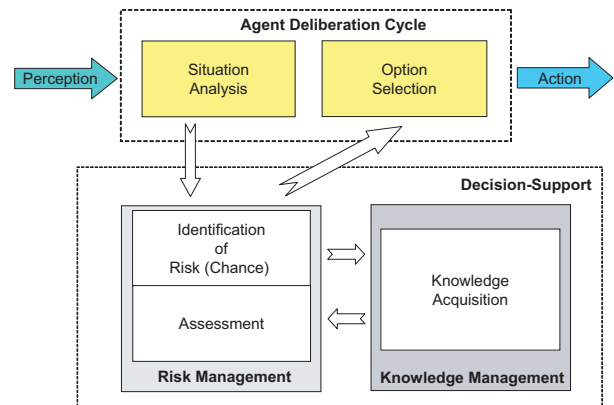


Figure 1: Interaction between agent deliberation, risk management, and knowledge management.

Our framework includes explicit *risk* and *knowledge* management, termed decision-support in the figure, which may work in an interweaved fashion to augment the deliberation cycle of the agent. Generally speaking, we use RM to identify and assess the risks associated with one or more options, and KM to acquire missing knowledge, for example, to improve risk assessment or to generate additional options. Our decision-support system can be integrated into any intelligent agent that utilizes some form of deliberation with separate option generation and selection phases.

Starting with the deliberation cycle at the top of Fig. 1, we assume that some perceptions are leading to a situation, where the agent has to decide on its next action. Before making a decision, the agent invokes RM to help with the assessment of the option(s). We are envisioning that all components have access to a common repository or knowledge base (not shown in the figure). For the remainder of the paper, we will use the term “beliefs” to refer to this knowledge base.

Agent-based or agent-mediated knowledge management is a rel-

*This research is funded by the German Research Foundation (DFG) as the Collaborative Research Centre 637 “Autonomous Cooperating Logistic Processes: A Paradigm Shift and its Limitations” (SFB 637). Additional Information can be found at <http://www.sfb637.uni-bremen.de/>.

atively young but currently very active field of research. Van Elst et al. [8] give a comprehensive overview. Risk assessment as scientific topic is basically known from management science, finance, environmentalism and health care. Therefore a number of methodologies for organizational risk identification and management can be found in the literature [2, 7]. Explicit, knowledge-based risk identification based solely on autonomously acquired data (in contrast to specific software-guided user input), i.e., a fully automated knowledge-based risk management system, has not yet been proposed in the considered literature.

2. RISK MANAGEMENT

Risk management is a continuous process that will trigger further deliberation as soon as a fact is added to the knowledge base, which makes the situation risky. Our concept of RM is heavily depending on knowledge. Therefore it can only function in close collaboration with a KM infrastructure. In the following we will describe the mechanisms of this collaboration and the core task of knowledge-based risk assessment [1].

The correlation between risk and knowledge management is at least threefold. First of all, knowledge of risk is one part of an agent's beliefs. Thus, it can be communicated by our approach of KM. Secondly, an agent can use its knowledge of the world to identify risks. From that point of view—the *knowledge-based risk identification*—knowledge is needed when the agent wants to reason about the possible risks it will face. And thirdly, the act of communicating knowledge is in itself a risk to the agent because it can fail in various ways.

The initial task and most important prerequisite for successful risk management is to identify risk and evaluate its potential consequence. In the situation analysis phase of an agent's deliberation cycle (see Fig. 1), incoming perceptions are integrated with the current beliefs (for details and a formal specification of this process we refer to recent work by Timm [6]). Risk identification will then search for incidents that may impact the execution of options. In the next step of risk assessment, the agent evaluates the evidences, i.e., beliefs, which are now tagged as risk relevant according to the degree of uncertainty it has about this evidence. Together with the gravity value, high uncertainty can trigger acquisition of additional knowledge.

In order to trigger the acquisition of new knowledge by KM, RM assigns a parameter *importance* to every item it requests.

To determine importance, the risk identification process interprets every perception as a belief that supports a given risk hypothesis. Every belief is associated with two probability values which denote the subjective confidence the agent has: *support* and *plausibility*. In this we follow the basic idea of the Dempster-Shafer theory of evidence (cf. [3] or [4]). Support for a hypothesis gives the amount of belief that directly supports a given hypothesis. Plausibility gives an upper bound on the belief that the hypothesis could possibly happen.

Based on this support values, the agent can express its need for new evidences. The *ignorance factor* denotes the agent's degree of uncertainty about facts needed to be able to soundly evaluate the risk in question. Together with a *gravity value* χ , these define the *importance* parameter for KM.

Based on a threshold for the ignorance factor, risk management decides whether the evidence that is already present in the beliefs is sufficiently crisp to assess the identified risk.

3. KNOWLEDGE MANAGEMENT

If the RM component identifies the need for additional informa-

tion, knowledge management is invoked as depicted in Fig. 1. Our approach to KM consists of three main components: **conceptual knowledge**, **roles**, and **parameters**. The conceptual knowledge is represented as an OWL¹ ontology. For the purpose of our logistic application domain, this ontology includes a representation of the transportation or production network, the basic types of agents and their properties (e.g., for a vehicle, its average and maximum speed, the types of routes in the network it can use, and its load capacity), and the properties of 'inactive' objects, such as highways, depots, etc.

In contrast to previous approaches to agent-based knowledge management [8], we do not presuppose a one-to-one correspondence between agents and KM functions, such as *providing* knowledge or *brokering* knowledge. In our approach these functions are implemented as *roles* [5]. A knowledge management role includes certain reasoning capabilities, a visibility function on an agent's beliefs, a deliberation pattern (i.e., a plan how to accomplish the KM task), and a communication behavior with interacting roles. The aim of KM roles is to provide a formal description of KM tasks that eases the development of agents and reduces computational complexity by means of a minimum set of processed knowledge and applied reasoning capabilities. One agent can assume different roles and may change them over time.

4. STATUS AND CONCLUSION

We have described our conceptual framework for enabling autonomous decision-making in agents. Our approach, which integrates risk and knowledge management, allows an agent to evaluate decisions/options based on the likelihood of certain beliefs that the agent uses as supporting evidence.

The approach described in this paper has the following three important benefits: (1) Our approach augments agent deliberation with sophisticated decision-making capabilities not found in current architectures. (2) By using RM to also support the acquisition of knowledge, our approach is better equipped to manage the highly dynamic, context-sensitive, and uncertain information needed to make autonomous decisions in realistic environments. This is of particular importance, since we do not presuppose benevolent behaviour. (3) Our role-based KM enables the distribution of knowledge and knowledge management functionality which eliminates the need for a centralized knowledge repository. On the other hand, it provides the necessary flexibility to allow existing KM infrastructure to co-exist with our approach.

5. REFERENCES

- [1] B. Bemeleit, M. Lorenz, J. Schumacher, and O. Herzog. Risk management for transportation of sensitive goods. In *Proceedings of the 10th. International Symposium on Logistic (10th ISL)*, Lisbon, Portugal, July 3 - 5 2005.
- [2] S. Figlewski and R. M. Levich, editors. *Risk management: the state of the art*. Kluwer Academic Press, 2002.
- [3] J. Y. Halpern. *Reasoning about uncertainty*. The MIT Press, 2003.
- [4] R. Kruse, E. Schwecke, and J. Heinsohn. *Uncertainty and Vagueness in Knowledge Based Systems*. Springer-Verlag, 1991.
- [5] H. Langer, J. D. Gehrke, J. Hammer, M. Lorenz, I. J. Timm, and O. Herzog. Emerging knowledge management in distributed environments. In *Workshop on Agent-Mediated Knowledge Management (AMKM 2005)*, pages 14–26, 2005.
- [6] I. J. Timm. *Dynamisches Konfliktmanagement als Verhaltenssteuerung Intelligenter Agenten*. PhD thesis, Universität Bremen, Bremen, Germany, Apr. 2004.
- [7] M. J. S. Timothy McDaniels, editor. *Risk analysis and society : an interdisciplinary characterization of the field*. Cambridge University Press, 2004.
- [8] L. van Elst, V. Dignum, and A. Abecker, editors. *Agent-Mediated Knowledge Management: International Symposium AMKM 2003, Stanford, CA, USA, March 24-26, Revised and Invited Papers*, volume 2926 of *Lecture Notes in Computer Science*. Springer-Verlag, 2004.

¹<http://www.w3.org/2004/OWL/>