



Deutsches
Forschungszentrum
für Künstliche
Intelligenz GmbH

**Technical
Memo**
TM-91-08

**Social and Psychological Commitments in
Multiagent Systems**

Munindar P. Singh

August 1991

**Deutsches Forschungszentrum für Künstliche Intelligenz
GmbH**

Postfach 20 80
D-6750 Kaiserslautern
Tel.: (+49 631) 205-3211/13
Fax: (+49 631) 205-3210

Stuhlsatzenhausweg 3
D-6600 Saarbrücken 11
Tel.: (+49 681) 302-5252
Fax: (+49 681) 302-5341

Deutsches Forschungszentrum für Künstliche Intelligenz

The German Research Center for Artificial Intelligence (Deutsches Forschungszentrum für Künstliche Intelligenz, DFKI) with sites in Kaiserslautern und Saarbrücken is a non-profit organization which was founded in 1988 by the shareholder companies ADV/Orga, AEG, IBM, Insiders, Fraunhofer Gesellschaft, GMD, Krupp-Atlas, Mannesmann-Kienzle, Philips, Siemens and Siemens-Nixdorf. Research projects conducted at the DFKI are funded by the German Ministry for Research and Technology, by the shareholder companies, or by other industrial contracts.

The DFKI conducts application-oriented basic research in the field of artificial intelligence and other related subfields of computer science. The overall goal is to construct *systems with technical knowledge and common sense* which - by using AI methods - implement a problem solution for a selected application area. Currently, there are the following research areas at the DFKI:

- Intelligent Engineering Systems
- Intelligent User Interfaces
- Intelligent Communication Networks
- Intelligent Cooperative Systems.

The DFKI strives at making its research results available to the scientific community. There exist many contacts to domestic and foreign research institutions, both in academy and industry. The DFKI hosts technology transfer workshops for shareholders and other interested groups in order to inform about the current state of research.

From its beginning, the DFKI has provided an attractive working environment for AI researchers from Germany and from all over the world. The goal is to have a staff of about 100 researchers at the end of the building-up phase.

Prof. Dr. Gerhard Barth
Director

Social and Psychological Commitments in Multiagent Systems

Munindar P. Singh

DFKI-TM-91-08

This work has been supported by a grant from The Federal Ministry for Research and Technology (FKZ ITW-8903 0).

© Deutsches Forschungszentrum für Künstliche Intelligenz 1991

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Deutsches Forschungszentrum für Künstliche Intelligenz, Kaiserslautern, Federal Republic of Germany; an acknowledgement of the authors and individual contributors to the work; all applicable portions of this copyright notice. Copying, reproducing, or republishing for any other purpose shall require a licence with payment of fee to Deutsches Forschungszentrum für Künstliche Intelligenz.

Abstract

Commitment is a central concept in Artificial Intelligence (AI). At least two kinds of commitment can be identified that have been used in AI—the internal or *psychological* and the external or *social*. While the former has been made explicit in AI theory, the latter has often been ignored (though it is given more importance in other disciplines, and indeed even in AI practice). The many roles of social commitment in AI are discussed. It is argued that while social and psychological commitments are related, they must not be conflated with each other. In particular, thinking directly in terms of social commitments helps us avoid the infelicities of traditional theories of group action and intention. In the full paper, a formalization of these concepts is also proposed that captures their desired properties and interrelationships.

1 Introduction

At least two kinds of commitment can be identified in work in Artificial Intelligence (AI). One is the familiar notion from Ethics and Distributed Computing where an agent is deemed to make a commitment to others to do certain actions (including actions such as issuing permissions) or to prevent certain conditions. The agent who makes such commitments may be held socially liable for not acting up to them. Another, quite different, sense of commitment has also been drawing attention recently. This is the kind of commitment that Bratman [1987, ch. 2] and Harman [1986, p. 94] invoke when they say that an agent is “committed” to his intentions. This is also the notion used in [Cohen and Levesque, 1990, p. 217] and studied in [Singh, 1991a]. As I will explain shortly, this sense of commitment also arises under the guise of the “epistemic entrenchment” of an agent’s beliefs as in [Gärdenfors, 1988] and [Harman, 1986]. I call the first kind of commitment *social* commitment or *S-commitment* and the second kind *psychological* commitment or *P-commitment*. These concepts are the subject of this paper.

The question: *can a group be the locus of believing and intending?* may be raised. The underlying assumption of this paper is that it can unequivocally be answered with a “Yes!” This seems justified by our commonsense intuitions about human and artificial multiagent systems. Groups such as teams, armies, nations and corporations may all be (and are in fact) profitably treated as being single, though potentially complex, agents. The notion of agency is intimately connected to ascriptions of belief and intention—if something can be considered an agent, it can and must be ascribed at least some relevant beliefs and intentions. That such ascriptions are legitimate has also been argued by McCarthy and Dennett: one only has to realize that groups are, at worst, artificial and complex physical systems [Dennett, 1987; McCarthy, 1979]. One can do better, however, since the “mental properties” of groups depend on the mental properties of their members and on how they are structured. Modern theories of intentions and beliefs, e.g., those of Harman, Bratman and Gärdenfors, all assert that believing or intending entails, or at least is intimately connected to, the notion of P-commitment. So the above question may be posed as: *can a group be the locus of P-commitment?*, and still be answered with a “Yes!” It is for this reason that in this paper, I take groups to be first-class agents. They are not only legitimate loci of P-commitments, but also of S-commitments, as explained below.

S-commitments arise when agents interact with other agents. They can be as complex as the social structures in which they arise. Just as P-commitments are needed in understanding simple agents, S-commitments are needed in understanding multiagent systems. Traditional formal theories give preeminence to P-commitment and do not seem to recognize the complexities of S-commitment that are possible in real systems—they allow S-commitments implicitly, but only in groups which are internally homogeneous, and even then are not quite correct. These theories do not quite make the mark even from a prescriptive or normative standpoint, because S-commitments are required when one wants to design complex multiagent systems: if we can design the appropriate kinds of S-commitments that come into play in a system, we can design the member agents of that system modularly—we only need to ensure that they generate and process their S-commitments properly. Indeed, if current theories are right, group action and intention are impossible! Note that since groups are first-class agents, they too can have S-commitments to other agents, who might themselves be groups.

The ideas of commitment (especially S-commitment) as explored in this paper underlie much research into distributed AI, though they are not usually formalized in the way proposed here; e.g., agents collaborating in design have to commit to various issues at different stages of the design to arrive at a consistent solution. The motivation for commitments as defined here is both descriptive and prescriptive. Firstly, commitments of both kinds are held by the intelligent systems that we encounter in real-life: humans, corporations and what have you. Secondly, commitments of both kinds can be motivated on grounds of individual and group rationality for limited systems and thus may be prescribed to them or incorporated in their design. Also, the problem of determining how commitments may be formalized so that their important properties can be captured is important—such formalizations usually involve complicated notions such as mutual belief and are, as I explain below, *a fortiori* unimplementable in real-life systems where communication is slow and unreliable. The approach proposed here is not only descriptively valid, but also prescriptively feasible.

As an explanation of my terminology, note that I use *psychological* to refer to anything within an agent, even if that agent is a group of possibly dumb agents, e.g., a termite colony in which the individual termites do not reflect on the colony's actions and do not have explicit symbolic representations of

them. I use *social* for anything arising between agents. Thus the title of this paper could have been “inter- and intra-agent commitments . . . ,” with the proviso that the intra-agent commitments of a group might rely on the inter-agent commitments of its members.

2 Commitments

It should be clear that whatever their interrelationships, P-commitments and S-commitments are, as concepts, quite different from each other:

1. An agent may be psychologically committed to a belief or an intention and yet not be S-committed to it.
2. An agent might be officially S-committed to an intention or belief, but not be psychologically committed to it; e.g., a bad cop might not intend to catch a robber friend of his.
3. An agent, e.g., Robinson Crusoe, can have psychological commitments, by himself, e.g., to trap a goat. But he needs another agent, e.g., Man Friday, to interact with to have S-commitments, e.g., to have dinner at sundown.

2.1 Psychological Commitments

P-commitments are comparatively straightforward and are only briefly discussed here. Intentions for future actions are an important concept in AI. Perhaps their salient property is that they involve a commitment on the part of agents. This view has been gaining ground in Philosophy and AI (e.g., see [Bratman, 1987, ch. 2], [Harman, 1986, p. 94] and [Cohen and Levesque, 1990, p. 217]). The idea here is that an agent who has an intention is in some way committed to it—not only does he intend to achieve the relevant condition right now, but would also intend to achieve it later, even as the circumstances changed, perhaps for the worse.

Agents can also be committed towards their beliefs, this commitment corresponding to the latter’s doxastic or epistemic entrenchment [Gärdenfors, 1988, ch. 3]. The more committed an agent is to a belief, the more disposed he will be to not reconsider it, even as he receives new evidence that potentially

challenges it or removes its justifications; e.g., if you firmly believe that all birds can fly, you might accept that penguins aren't birds, rather than that some birds cannot fly.

P-commitment entails that the agent continue with a belief that might no longer be justified and an intention that might no longer be feasible or serve the agent's ultimate goals. Thus there is a certain amount of irrationality built into this concept. But for limited agents, commitments can be reconciled with rationality. Limited agents cannot be expected to reason at every moment from "first principles"—this kind of reasoning can be simply too expensive to carry out repeatedly. P-commitments help the agent consider only a few issues, which he can hope to resolve in the limited time he has available. The reader may consult [Singh, 1991a] for the details of one possible mechanism by which P-commitments might help. I will just take it as granted here that it is a feature of good design of agents that they are able to manage despite its limitations.

2.2 Social Commitments

This kind of commitment arises often in distributed AI, e.g., when the agent promises (to another) to do a certain action. The S-commitment of an agent (to another) in a multiagent system to achieve a certain goal is essential to coordination; e.g., if one agent promises another to be at a rendezvous spot, then the other can take this for granted and be there himself; or if one agent promises to lift one end of a piano, the other can lift the other end, and jointly they can lift the piano, thereby achieving something they could not have achieved singly. In general, I take S-commitment as the commitment of an agent to another agent. The second agent may possibly be a group that includes the first agent—this definition is thus sufficiently powerful to account for what one might call "team spirit," or the S-commitments agents due to it. It is important to note that commitments in classical Distributed Computing are irrevocable—this makes much of this research inapplicable to AI, where such rigidity would be undesirable.

Whether or not S-commitments may be reduced to P-commitments in principle, they are at least in practice an important concept for AI. It would be impossible to design or understand sufficiently complex systems of autonomous agents without invoking the concept of S-commitment. The agents' S-commitments would give us a powerful abstraction with which to view their

interactions. Thus, just as psychological concepts are needed to understand complex agents, social concepts are needed to understand systems of agents. And S-commitment is a social concept *par excellence*.

1 S-commitments as interfaces:

S-commitments are a way of specifying interfaces between intelligent agents in multiagent systems. The most interesting interactions among agents involve communications, whether implicit or explicit. Communications can be described fruitfully using social commitment. For example, promises bring into effect a social commitment on part of the speaker to the hearer; directives presuppose an S-commitment on part of the hearer to do as told; assertives S-commit the speaker to the statement expressed (even if the speaker is insincere, in which case he would not be P-committed to it); permissions make the speaker S-committed to allowing the relevant condition to hold; prohibitions presuppose an S-commitment by the hearer to preventing the relevant condition from holding in the relevant time-frame. Interaction protocols between different agents may be defined so that the interacting agents have the relevant commitments. This generalizes the idea of [Singh, 1991b] that participating agents have the requisite intentions and know-how. This idea can also help differentiate communication from merely physical interactions such as resource conflicts. Even though an agent might obtain some information from another because of physical actions such as resource conflicts, these interactions differ from the strictly communicative ones in lacking the properties discussed above; e.g., the agents competing for a resource are not S-committed to requiring it.

A practical consequence of having S-commitments and treating them as interfaces is that one can design agents independently of each other. We would just need to ensure that their S-commitments would mesh in properly with others' S-commitments when they are combined into a distributed system.

2 S-commitments and individual rationality:

An agent can *bargain* with other agents about the adoption of some S-commitments. If an agent agrees to do something that another agent intends he do, he might be able to get something in return from the

other agent. If the first agent just intended to do something without S-committing to it (to the second agent), he would have no grounds for bargaining with him; e.g., Robinson Crusoe can get Man Friday to catch a grouse for them only by S-committing to milk the goat for them. If he just planned to milk the goat for them, but did not S-commit to doing so, he would not be able to use that as a basis for convincing Man Friday to perform his share of the chores.

Another potential benefit to an agent is that S-commitments can make it easier for the group as a whole to achieve its ends, and this might have some positive utility for him. For example, Crusoe and Friday might need to work as a team to defend their island, a task they would be able to do better if each S-commits to watching one side of it. Their success as a team is beneficial to each of them individually.

3 S-commitments and group rationality:

The role of S-commitments in group rationality is akin to that of P-commitments in individual rationality. Even when they are not in the immediate (i.e., local) interest of an agent, S-commitments might be a good idea from the systems point of view. When the agents in a system make and keep S-commitments, the system can perform better than it might have done if they each had worked selfishly. This too depends on specific systems, but there are numerous real-life systems where S-commitments alone make it possible for any kind of success; e.g., the famous prisoners' dilemma paradox arises when the individually preferred actions of two agents (namely to tell on each other) lead to a worse pay-off for each of them than the individually less preferred actions of protecting each other. The latter action can be done reasonably only if the agents are S-committed to each other to not aid the police.

4 S-commitments and representation:

S-commitments may be implicit; i.e., they need not be explicitly symbolically represented by the agents, but could instead be derived from their social architectures. Of course, they *could* be symbolically represented by *some* smart introspective agents. Often the relevant agents would be treated only indexically, i.e., relative to the agent's own situation. For example, at a stop sign at a street crossing, the other agent

would be “the driver to my right,” rather than Bill, the town councilman. The S-commitment itself to let him go (and the S-commitment on his part to get out of my way) are both a matter of training, here, of learning how to drive. S-commitments may also be generated by the *social roles* of agents; e.g., someone, merely by being a policeman in uniform, is S-committed to chase after a criminal. His not doing so would constitute a dereliction of duty.

5 Commitments and coordination:

In traditional construals of group action, the agents in the group are required to have mutual beliefs about the actions performed by the group [Levesque *et al.*, 1990; Grosz and Sidner, 1988]. Here it is proposed that no such mutual beliefs are necessary. When two agents make social commitments to each other, they are already in a position to perform their joint action. Mutual beliefs are difficult to achieve in realistic scenarios and are highly unstable even when achieved—the slightest suspicion by one of the participants that the other no longer has the relevant belief (to any arbitrary nesting) is grounds for the failure of the mutual belief; e.g., if one agent comes to doubt (for a large natural number, n) that the other believes that he believes ... (n times) the given proposition, their mutual belief no longer exists, even if this doubt is unfounded. This is a fundamental problem for theories that rely on mutual beliefs to explain or define coordination. No number of epicycles added to such a theory can repair this problem—as long as there are mutual beliefs, they will be unstable. Indeed, the requirement of mutual belief may be taken as a purported proof that joint action is impossible (e.g., as in [Fischer and Immerman, 1986; Halpern and Moses, 1987]).

A far simpler definition for coordination can be motivated: the involved agents merely have to have the appropriate S-commitments to one another to act in an appropriate manner. Two agents are jointly committed to achieving a condition, p , iff each of them is S-committed to the other to achieve a (possibly different) condition such that the achievement of those two conditions would entail the achievement of p . The discussion below on revoking commitments applies here if an agent needs to drop out of some joint commitment.

2.2.1 Revoking Social Commitments

While P-commitments can be revoked at will (at least in principle), S-commitments cannot be revoked till the agent who is committed (called the *committer* in the sequel) clarifies this to the other agent (called the *committee* in the sequel). Ideally, an S-commitment would be terminated when the requirements it imposed were met (i.e., the committee is satisfied), or when the committee released the commiter. These roles make the committee very important.

An agent might also need to give up his S-commitment if the circumstances change, e.g., if his ultimate goals change. As a matter of descriptive fact for humans, the agent might then simply drop his S-commitment and quit acting for it. This kind of unreliable behavior can lead to a number of complications, e.g., as when someone does not show up for an appointment. Prescriptively speaking, we would like to prevent such behavior, at least to a first approximation (we might not always want to prevent it because it might sometimes be acceptable—e.g., someone who foolishly agrees to rob a bank, might be forgiven for not showing up). But in the more mundane cases it helps a system if its members do not drop their S-commitments arbitrarily.

We should require that to revoke a S-commitment, the commiter must first get an acknowledgement from the committee. All that we need is that when a commitment is revoked, the commiter knows that the committee knows that. Mutual beliefs are not needed since once the committee believes the commitment is being given up, he need no longer act as if it held. It would be unwise for the committee to act as if the commitment persisted even though he knows it does not (merely on the hope that the commiter might persist with it a little longer yet). Not requiring mutual beliefs to be established makes for a considerable simplification, by avoiding the shortcomings of traditional theories as discussed in §2.2.

References

- [Bratman, 1987] Bratman, Michael E. 1987. *Intention, Plans and Practical Reason*. Harvard University Press, Cambridge, MA.
- [Cohen and Levesque, 1990] Cohen, Philip R. and Levesque, Hector J. 1990. Intention is choice with commitment. *Artificial Intelligence* 42:213–261.

- [Dennett, 1987] Dennett, Daniel C. 1987. *The Intentional Stance*. MIT Press, Cambridge, MA.
- [Fischer and Immerman, 1986] Fischer, Michael J. and Immerman, Neil 1986. Foundations of knowledge for distributed systems. In Halpern, Joseph Y., editor 1986, *Theoretical Aspects of Reasoning About Knowledge*. 171-185.
- [Gärdenfors, 1988] Gärdenfors, Peter 1988. *Knowledge in Flux*. MIT Press, Cambridge, MA.
- [Grosz and Sidner, 1988] Grosz, Barbara and Sidner, Candace 1988. Plans for discourse. In Cohen, P.; Morgan, J.; and Pollack, M., editors 1988, *SDF Benchmark Series: Intentions in Communication*. MIT Press, Cambridge, MA.
- [Halpern and Moses, 1987] Halpern, Joseph Y. and Moses, Yoram O. 1987. Knowledge and common knowledge in a distributed environment (revised version). Technical Report RJ 4421, IBM.
- [Harman, 1986] Harman, Gilbert 1986. *Change in View*. MIT Press, Cambridge, MA.
- [Levesque *et al.*, 1990] Levesque, H. J.; Cohen, P. R.; and Nunes, J. T. 1990. On acting together. In *AAAI-90*.
- [McCarthy, 1979] McCarthy, John 1979. Ascribing mental qualities to machines. In Ringle, Martin, editor 1979, *Philosophical Perspectives in Artificial Intelligence*. Harvester Press. Page nos. from a revised version, issued as a report in 1987.
- [Singh, 1991a] Singh, Munindar P. 1991a. Intentions, commitments and rationality. In *13th Annual Conference of the Cognitive Science Society*.
- [Singh, 1991b] Singh, Munindar P. 1991b. Towards a formal theory of communication for multiagent systems. In *International Joint Conference on Artificial Intelligence*.



Deutsches
Forschungszentrum
für Künstliche
Intelligenz GmbH

DFKI
-Bibliothek-
PF 2080
6750 Kaiserslautern
FRG

DFKI Publikationen

Die folgenden DFKI Veröffentlichungen oder die aktuelle Liste von erhältlichen Publikationen können bezogen werden von der oben angegebenen Adresse.

Die Berichte werden, wenn nicht anders gekennzeichnet, kostenlos abgegeben.

DFKI Publications

The following DFKI publications or the list of currently available publications can be ordered from the above address.

The reports are distributed free of charge except if otherwise indicated.

DFKI Research Reports

RR-90-01

Franz Baader: Terminological Cycles in KL-ONE-based Knowledge Representation Languages
33 pages

RR-90-02

Hans-Jürgen Bürckert: A Resolution Principle for Clauses with Constraints
25 pages

RR-90-03

Andreas Dengel, Nelson M. Mattos: Integration of Document Representation, Processing and Management
18 pages

RR-90-04

Bernhard Hollunder, Werner Nutt: Subsumption Algorithms for Concept Languages
34 pages

RR-90-05

Franz Baader: A Formal Definition for the Expressive Power of Knowledge Representation Languages
22 pages

RR-90-06

Bernhard Hollunder: Hybrid Inferences in KL-ONE-based Knowledge Representation Systems
21 pages

RR-90-07

Elisabeth André, Thomas Rist: Wissensbasierte Informationspräsentation:
Zwei Beiträge zum Fachgespräch Graphik und KI:
1. Ein planbasierter Ansatz zur Synthese illustrierter Dokumente
2. Wissensbasierte Perspektivenwahl für die automatische Erzeugung von 3D-Objektdarstellungen
24 pages

RR-90-08

Andreas Dengel: A Step Towards Understanding Paper Documents
25 pages

RR-90-09

Susanne Biundo: Plan Generation Using a Method of Deductive Program Synthesis
17 pages

RR-90-10

Franz Baader, Hans-Jürgen Bürckert, Bernhard Hollunder, Werner Nutt, Jörg H. Siekmann: Concept Logics
26 pages

RR-90-11

Elisabeth André, Thomas Rist: Towards a Plan-Based Synthesis of Illustrated Documents
14 pages

RR-90-12

Harold Boley: Declarative Operations on Nets
43 pages

RR-90-13

Franz Baader: Augmenting Concept Languages by Transitive Closure of Roles: An Alternative to Terminological Cycles
40 pages

RR-90-14

Franz Schmalhofer, Otto Kühn, Gabriele Schmidt: Integrated Knowledge Acquisition from Text, Previously Solved Cases, and Expert Memories
20 pages

RR-90-15

Harald Trost: The Application of Two-level Morphology to Non-concatenative German Morphology
13 pages

RR-90-16

Franz Baader, Werner Nutt: Adding Homomorphisms to Commutative/Monoidal Theories, or: How Algebra Can Help in Equational Unification
25 pages

RR-90-17

Stephan Busemann: Generalisierte Phasenstrukturgrammatiken und ihre Verwendung zur maschinellen Sprachverarbeitung
114 Seiten

RR-91-01

Franz Baader, Hans-Jürgen Bürckert, Bernhard Nebel, Werner Nutt, Gert Smolka: On the Expressivity of Feature Logics with Negation, Functional Uncertainty, and Sort Equations
20 pages

RR-91-02

Francesco Donini, Bernhard Hollunder, Maurizio Lenzerini, Alberto Marchetti Spaccamela, Daniele Nardi, Werner Nutt: The Complexity of Existential Quantification in Concept Languages
22 pages

RR-91-03

B.Hollunder, Franz Baader: Qualifying Number Restrictions in Concept Languages
34 pages

RR-91-04

Harald Trost: X2MORF: A Morphological Component Based on Augmented Two-Level Morphology
19 pages

RR-91-05

Wolfgang Wahlster, Elisabeth André, Winfried Graf, Thomas Rist: Designing Illustrated Texts: How Language Production is Influenced by Graphics Generation.
17 pages

RR-91-06

Elisabeth André, Thomas Rist: Synthesizing Illustrated Documents A Plan-Based Approach
11 pages

RR-91-07

Günter Neumann, Wolfgang Finkler: A Head-Driven Approach to Incremental and Parallel Generation of Syntactic Structures
13 pages

RR-91-08

Wolfgang Wahlster, Elisabeth André, Som Bandyopadhyay, Winfried Graf, Thomas Rist: WIP: The Coordinated Generation of Multimodal Presentations from a Common Representation
23 pages

RR-91-09

Hans-Jürgen Bürckert, Jürgen Müller, Achim Schupeta: RATMAN and its Relation to Other Multi-Agent Testbeds
31 pages

RR-91-10

Franz Baader, Philipp Hanschke: A Scheme for Integrating Concrete Domains into Concept Languages
31 pages

RR-91-11

Bernhard Nebel: Belief Revision and Default Reasoning: Syntax-Based Approaches
37 pages

RR-91-12

J.Mark Gawron, John Nerbonne, Stanley Peters: The Absorption Principle and E-Type Anaphora
33 pages

RR-91-13

Gert Smolka: Residuation and Guarded Rules for Constraint Logic Programming
17 pages

RR-91-14

Peter Breuer, Jürgen Müller: A Two Level Representation for Spatial Relations, Part I
27 pages

RR-91-15

Bernhard Nebel, Gert Smolka: Attributive Description Formalisms ... and the Rest of the World
20 pages

RR-91-16

Stephan Busemann: Using Pattern-Action Rules for the Generation of GPSG Structures from Separate Semantic Representations
18 pages

RR-91-17

Andreas Dengel & Nelson M. Mattos: The Use of Abstraction Concepts for Representing and Structuring Documents
17 pages

RR-91-19

Munindar P. Singh: On the Commitments and Precommitments of Limited Agents
15 pages

RR-91-20

Christoph Klauck, Ansgar Bernardi, Ralf Legleitner: FEAT-Rep: Representing Features in CAD/CAM
48 pages

RR-91-22

Andreas Dengel: Self-Adapting Structuring and Representation of Space
27 pages

RR-91-23

Michael Richter, Ansgar Bernardi, Christoph Klauck, Ralf Legleitner: Akquisition und Repräsentation von technischem Wissen für Planungsaufgaben im Bereich der Fertigungstechnik
24 Seiten

RR-91-24

Jochen Heinsohn: A Hybrid Approach for Modeling Uncertainty in Terminological Logics
22 pages

RR-91-25

Karin Harbusch, Wolfgang Finkler, Anne Schauder: Incremental Syntax Generation with Tree Adjoining Grammars
16 pages

RR-91-26

M. Bauer, S. Biundo, D. Dengler, M. Hecking, J. Koehler, G. Merziger: Integrated Plan Generation and Recognition - A Logic-Based Approach -
17 pages

RR-91-27

A. Bernardi, H. Boley, Ph. Hanschke, K. Hinkelmann, Ch. Klauck, O. Kühn, R. Legleitner, M. Meyer, M. M. Richter, F. Schmalhofer, G. Schmidt, W. Sommer: ARC-TEC: Acquisition, Representation and Compilation of Technical Knowledge
18 pages

RR-91-30

Dan Flickinger, John Nerbonne: Inheritance and Complementation: A Case Study of Easy Adjectives and Related Nouns
39pages

DFKI Technical Memos
TM-89-01

Susan Holbach-Weber: Connectionist Models and Figurative Speech
27 pages

TM-90-01

Som Bandyopadhyay: Towards an Understanding of Coherence in Multimodal Discourse
18 pages

TM-90-02

Jay C. Weber: The Myth of Domain-Independent Persistence
18 pages

TM-90-03

Franz Baader, Bernhard Hollunder: KRIS: Knowledge Representation and Inference System -System Description-
15 pages

TM-90-04

Franz Baader, Hans-Jürgen Bürckert, Jochen Heinsohn, Bernhard Hollunder, Jürgen Müller, Bernhard Nebel, Werner Nutt, Hans-Jürgen Profitlich: Terminological Knowledge Representation: A Proposal for a Terminological Logic
7 pages

TM-91-01

Jana Köhler: Approaches to the Reuse of Plan Schemata in Planning Formalisms
52 pages

TM-91-02

Knut Hinkelmann: Bidirectional Reasoning of Horn Clause Programs: Transformation and Compilation
20 pages

TM-91-03

Otto Kühn, Marc Linster, Gabriele Schmidt: Clamping, COKAM, KADS, and OMOS: The Construction and Operationalization of a KADS Conceptual Model
20 pages

TM-91-04

Harold Boley: A sampler of Relational/Functional Definitions
12 pages

TM-91-05

Jay C. Weber, Andreas Dengel, Rainer Bleisinger: Theoretical Consideration of Goal Recognition Aspects for Understanding Information in Business Letters
10 pages

TM-91-08

Munindar P. Singh: Social and Psychological Commitments in Multiagent Systems
11 pages

TM-91-09

Munindar P. Singh: On the Semantics of Protocols Among Distributed Intelligent Agents
18 pages

TM-91-10

Béla Buschauer, Peter Poller, Anne Schauder, Karin Harbusch: Tree Adjoining Grammars mit Unifikation
149 pages

TM-91-11

Peter Wazinski: Generating Spatial Descriptions for Cross-modal References
21 pages

DFKI Documents**D-89-01**

Michael H. Malburg, Rainer Bleisinger:
HYPERBIS: ein betriebliches Hypermedia-
Informationssystem
43 Seiten

D-90-01

DFKI Wissenschaftlich-Technischer Jahresbericht
1989
45 pages

D-90-02

Georg Seul: Logisches Programmieren mit Feature-
-Typen
107 Seiten

D-90-03

*Ansgar Bernardi, Christoph Klauck, Ralf
Legleitner*: Abschlußbericht des Arbeitspaketes
PROD
36 Seiten

D-90-04

*Ansgar Bernardi, Christoph Klauck, Ralf
Legleitner*: STEP: Überblick über eine zukünftige
Schnittstelle zum Produktdatenaustausch
69 Seiten

D-90-05

*Ansgar Bernardi, Christoph Klauck, Ralf
Legleitner*: Formalismus zur Repräsentation von
Geo-metrie- und Technologieinformationen als Teil
eines Wissensbasierten Produktmodells
66 Seiten

D-90-06

Andreas Becker: The Window Tool Kit
66 Seiten

D-91-01

Werner Stein, Michael Sintek: Relfun/X - An
Experimental Prolog Implementation of Relfun
48 pages

D-91-03

*Harold Boley, Klaus Elsbernd, Hans-Günther Hein,
Thomas Krause*: RFM Manual: Compiling
RELFUN into the Relational/Functional Machine
43 pages

D-91-04

DFKI Wissenschaftlich-Technischer Jahresbericht
1990
93 Seiten

D-91-06

Gerd Kamp: Entwurf, vergleichende Beschreibung
und Integration eines Arbeitsplanerstellungssystems
für Drehteile
130 Seiten

D-91-07

Ansgar Bernardi, Christoph Klauck, Ralf Legleitner
TEC-REP: Repräsentation von Geometrie- und
Technologieinformationen
70 Seiten

D-91-08

Thomas Krause: Globale Datenflußanalyse und
horizontale Compilation der relational-funktionalen
Sprache RELFUN
137 pages

D-91-09

David Powers and Lary Reeker (Eds.):
Proceedings MLNLO'91 - Machine Learning of
Natural Language and Ontology
211 pages
Note: This document is available only for a
nominal charge of 25 DM (or 15 US-\$).

D-91-10

Donald R. Steiner, Jürgen Müller (Eds.)
MAAMAW'91: Pre-Proceedings of the 3rd
European Workshop on „Modeling Autonomous
Agents and Multi-Agent Worlds“
246 pages
Note: This document is available only for a
nominal charge of 25 DM (or 15 US-\$).

D-91-11

Thilo C. Horstmann: Distributed Truth Maintenance
61 pages

D-91-12

Bernd Bachmann:
Hiera_{Con} - a Knowledge Representation System
with Typed Hierarchies and Constraints
75 pages

D-91-13

International Workshop on Terminological Logics
*Organizers: Bernhard Nebel, Christof Peltason, Kai
von Luck*
131 pages

D-91-14

*Erich Achilles, Bernhard Hollunder, Armin Laux,
Jörg-Peter Mohren*: KRJS: Knowledge
Representation and Inference System
- Benutzerhandbuch -
28 Seiten

D-01-00
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-01
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-02
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-03
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-04
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-05
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-06
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-07
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-08
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-09
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-10
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-11
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-12
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-13
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-14
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-15
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-16
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-17
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-18
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-19
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-20
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-21
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-22
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-23
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-24
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-25
 Die Entwicklung der ...
 12. Jahrgang
 1987


D-01-26
 Die Entwicklung der ...
 12. Jahrgang
 1987


D-01-27
 Die Entwicklung der ...
 12. Jahrgang
 1987


D-01-28
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-29
 Die Entwicklung der ...
 12. Jahrgang
 1987

D-01-30
 Die Entwicklung der ...
 12. Jahrgang
 1987







Social and Psychological Commitments in Multiagent Systems

Munindar P. Singh

TM-91-08

Technical Memo