Curriculum vitae for Erik Sandewall

Latest update 2014-03-06

Outline information

Current affiliation

Professor of computer science (retired) at the Department of Computer and Information Science at Linköping University, Sweden.

Milestones in professional life

- Year of birth: 1945.
- Finish high school, enter Uppsala University, Sweden: 1963.
- Bachelor degree (fil.kand.-examen): 1964.
- Graduate student at Stanford University, USA: 1966–67.
- Ph.D. from Uppsala University: 1969.
- Docent at Uppsala University: 1969–1974.
- Visiting Associate Professor at MIT AI Lab, USA: 1974–75.
- Professor of Computer Science at Linköping University: 1975–present.
- Visiting Researcher at LAAS laboratory, Toulouse, France: 1993–94.
- Visiting Associate Professor (part time) at the Royal Institute of Technology, Stockholm, Sweden: 2007–2011.
- Retired: 2012.
- Plan to pursue research at least until the end of 2015.

Memberships and honors

International/Europe/USA

Member of Academia Europaea.
Fellow of the European Coordinating Committee for Artificial Intelligence.
Fellow of the Association for the Advancement of Artificial Intelligence.
Recipient of the ECCAI Distinguished service award (2012).

Sweden

Member of the Royal Swedish Academy of Sciences (KVA).
Member of the Royal Swedish Academy of Engineering Sciences (IVA).

Royal Swedish Academy of Sciences, Award to Outstanding Young Researcher, 1969.
Royal Swedish Academy of Engineering Sciences, Chester Carlson prize, 1985.

Member of the Technical-Scientific Council of the Swedish Agency for Non-proliferation and Export Controls.

France

Knight of the Legion of Honour.
Doctor Honoris Causa at Université Paul Sabatier, Toulouse.
President of the French-Swedish Research Society (AFSR).

Germany

Fellow of the Deutsches Forschungsinstitut für Künstliche Intelligenz (DFKI).

Morocco

Associate member of the Hassan II Academy of Science and Technology.

Past leadership positions

Member twice of *Forskningsberedningen* – advisory group on research matters to the Swedish government.

Prorector of Linköping University, circa 1993-2001.

Personal life

Residence: in the city of Linköping, Sweden.

Married, have two sons.

E-mail address: `erik.sandewall@liu.se`

Research activities and contributions

Research training

As shown in the previous section, I obtained my Ph.D. in 1969 from Uppsala university in Sweden. While working for the Ph.D. I had spent one year as a graduate student at the Stanford Artificial Intelligence Laboratory (1966–67), with Professor John McCarthy as my principal advisor. Professor Heinz-Otto Kreiss who is an authority in the field of Numerical analysis acted as my advisor in Uppsala.

Additional, shorter visits at Stanford university and a one-year stay at the Massachusetts Institute of Technology during my docent period were also important for introducing me to doing research in the field of Artificial intelligence.

Particular Basic Research Contributions

The following is an outline of those research topics that I have addressed during my career and that I want to mention in particular. Numbers in brackets are references to the list of my selected research publications that is included as the final section of this document.

Nonmonotonic Reasoning

My article (35) was the first to propose dealing with AI-style nonmonotonicity using an extension to standard first-order logic, and to identify the phenomenon of multiple extensions that is a consequence of that extension. Additional contributions to this topic are (9,51,28).

Reasoning about Actions

From 1985 and ten years onwards my primary interest was formal methods for reasoning about actions and change. The first step was the formulation and use of state-transition semantics for this purpose (47,49). After this I identfied two significant techniques for such reasoning, namely the methods of *occlusion* and of *filtering* (53,52) which have then been adopted by several other authors although sometimes under different names.

I also showed how the method of chronological minimization of change, which is due to Yoav Shoham, can be nicely generalized to chronological minimization of discontinuities, for the purpose of reasoning about actions and change over continuous time and with continuous-valued fluents (53)

My research monograph *Features and Fluents* (1994) and articles (59,58,14,77) introduced a systematic framework for analyzing the range of applicability for methods of reasoning about actions and change. These publications introduced trajectory semantics as a generalization of state-transition semantics, and used it for the analysis of both previously published methods, the ones mentioned

above, and a number of others. It also presented upper and lower bounds (often identical) for the range of applicability of a number of such methods, both existing ones in the literature, and several new ones. One of them, namely the PMON method, is used for the TAL logic and system that have been developed and are being used very successfully by my student Patrick Doherty and by his students in turn.

For a similar analysis of ramification methods I introduced transition cascade semantics and used it to obtain range of applicability results for approaches to ramification that are based on minimization of change (62,61).

These semantics were used for the representation of specific aspects of cognitive agent systems (16,18,63,65) and further articles discussed how the theoretical results could be applied in actual software system architectures (55,67,80).

The work in this area was mostly done in the 1990's.

Partial Evaluation

My group was one of several groups that independently developed and used the method of partial evaluation for use in AI and other software technology. My particular interest in it was as a tool for the efficient use of special-purpose languages, which are of course important in AI. Our initial work on partial evaluation (7) was followed by several PhD degrees in our group (Haraldsson, Emanuelsson, Komorowski) and was also picked up by the group of J. Fischer-Nielsen in Copenhagen. This work was done in the 1970's. A later article (20) puts this early work into perspective.

Workflow

My group developed what we called "information flow" in the late 1970's, meaning what is today popular under the name of workflow. Our work eg. in (40,44) was concurrent with, and independent of the work at MIT at that time (by Tom Malone) but to my knowledge there was no other work on this topic at that time or before.

Three-level Architecture

The article (13) and related articles by some of my students are among the very first publications on the topic of three-level architectures for autonomous systems, which is now a very widely used approach.

Defeasible Inheritance

For a long time I have been trying to find a semantically wellfounded approach to defeasible inheritance, in a way that could parallel the analysis of methods of reasoning about actions and change as described below. One earlier article

on this topic is (10). I believe that I have now found the definite approach to this problem, as reported in the AIJ article (27).

Large Applied Research Projects

Besides a considerable number of small to medium size projects over the years, I have also been the coordinator on the European level for the Artificial intelligence part of the Prometheus project (54). I was also the main responsible for the WITAS project (66) which was financed by a major grant from the Wallenberg foundation.

Strategic Efforts and Contributions

The following are efforts and contributions that I consider to be at least as important as those listed above, but which do not have the character of specific results on specific research problems. This also means that they are not fully reflected in conventional journal and conference articles.

Bridging AI and Software Technology

Before the 1980's, research in Artificial Intelligence took an active interest in the software technology for programming languages and interactive programming systems, in particular through the Lisp and Prolog languages. This is no longer the case, and AI research seems to be content with generally accepted technologies such as those based on Java and on XML. I believe that this is a strategic mistake for the field, and that AI should resume an interest in software technology.

This question of software technology for AI has been a recurrent interest for me since I started in the field, as published for example in (31,5,38,8,56,12). During the last fifteen years I have further developed this earlier work into an approach based on the concept of *software individuals,* ie. a kind of persistent, self-aware, autonomous intelligent agents (22,70). The *Leonardo system* is an experimental implementation of these concepts (81,73,74). The work on the Leonardo system started in 2005.

The IT4 Research Program and Plan-based Systems

The research in the groups that I have led have of course been heavily dependent on funding from outside our own university. It would not be possible to list all those grants here, and only a few will be mentioned. Besides the grants for our participation in the Prometheus project and the grant for the WITAS project, which were mentioned above, it is also appropriate to mention our project that was entitled *Plan-based systems* and that was part of the Swedish IT4 research program in the 1980's. This project focussed on one important aspect of AI systems, namely, their ability to plan their own behavior, execute these plans, but also to modify the plans if conditions require so.

The concept of Plan-based systems served as the framework for my own research on reasoning about actions which was described above. However, the IT4 grant for Plan-based systems also provided the initial resources for the development of two lines of research in Linkping: Patrick Doherty's work that led up to TALPlanner, and Christer Bckstrms work (continued by Peter Jonsson) on complexity analysis of planning algorithms. Both of these lines of research have made very significant contributions to the research in the field.

The IT4 research program also funded our participation in the Prometheus project, although as a separate grant.

Open Peer Review and Open Access

In 1997 I started the *Electronic Transactions on Artificial Intelligence,* (ETAI) located at `http://www.etaij.org/`, as a journal that would use a novel approach to peer review, as described in (17). It was the first journal to use this approach, to the best of my knowledge. The ETAI used an open-access publication method, but this was a means rather than an end in itself.

The ETAI approach opened up some issues concerning the principles for defining and certifying electronic publications, and the formation of a study group that would report to the International Association of STM Publishers. The report from the study group, where I was one of the members, has been published as (21).

The possibility of open review did not however get very much attention or following during the first ten years, but more recently there has been a growing interest in such alternative peer review methods, as witnessed by special issues in *Nature* and in the journal *Frontiers in Computational Neuroscience* that were devoted to this topic. Both of them contained articles on the ETAI experience, (23,29).

Linköping University Electronic Press

In order to implement the open-review policy of the ETAI, it was necessary to have a means of open-access publication, and for two purposes: for the publication of contributed articles while they were in the review period, and also for the publication of the journal as such, containing those articles that had been accepted following review. In 1996 I therefore proposed the creation of the Linköping University Electronic Press as a separate entity within the university, and I led its work until year 2007. Its original mission was to publish a few electronic journals as well as technical reports from some of our departments. It has gradually broadened its scope, and it now serves the entire university.

Research Methodology in Artificial Intelligence

Much work in Artificial Intelligence has the character of explorative system building - software systems are built, not for the purpose of using them in production mode, nor in order to make performance measurements on particular

algorithms, but instead in order to gain experience with a particular software architecture or other similar design principle. I propose that there is a dire need to clarify the methodology for such research, in particular, what can be claimed as a result from such a project, and what is required as evidence for such claimed results. The website `http://www.ida.liu.se/ext/morador/` contains an article and ensuing debate on this topic. This is a thesis that I have proposed repeatedly since the early 1970's. The discussion proceeds on a person-to-person level but it rarely arrives into publications.

Comprehensive Textbook on Knowledgebased Artificial Intelligence

The field of Artificial Intelligence has become increasingly fragmented, and one consequence of this is that current AI textbooks tend to describe the field through a catalog of different subfields, each with its own terminology and notation. I consider this to be a strategic problem for the field, and my contribution towards a solution is to develop a textbook that presents the major areas of knowledgebased AI in a systematic and unified fashion. This will be an open-access textbook (Open Educational Resource) that will come in several parts. The webpage for this project is at `http://www.ida.liu.se/ext/aica/`. The textbook shares the same conceptual framework and the same notation as the Leonardo system that was described above.

Research Community Activities

The present section lists major activities in support of the research community in Artificial intelligence.

Major A.I. Conferences

I was a member of the conference committee for the first (1969) International Joint Conference on Artificial Intelligence, IJCAI, with the role of European coordinator. I then served in the group of trustees of this conference series during 1973-1981, and in particular I was elected general chair for the 1975 IJCAI conference that was held in Tbilisi in the then Soviet Union. This turned out to be a many-faceted task due to the cold-war political turbulence that had an impact on the conference.

I also led the bid for having the 1990 European Conference on Artificial Intelligence (ECAI) in Stockholm, and as the bid was successful I had the role of local arrangements chair.

Furthermore I was co-program-chair (with Richard Fikes) for the 1991 Conference on Knowledge Representation (KR), and subsequently general chair for the 1994 KR conference.

Carl-Gustaf Jansson and I were co-program-chairs of the 1993 Scandinavian Conference on Artificial Intelligence.

When the 1999 IJCAI conference was held in Stockholm it was co-hosted by the Scandinavian countries, and at that point I was in charge of the Scandinavian level coordination.

Specialized conferences

I have been a co-editor of the following books with collections of articles at a specialized international conferences.

- With Pier-Paolo Degano, *Integrated Interactive Computing Systems,* (North-Holland, 1983).
- With David Barstow and Howard Shrobe, *Interactive Programming Environments,* (McGraw-Hill, 1984). – Both this workshop and the preceding one addressed the interaction between AI and software technology, in particular with respect to the style of interactive software development that has been pioneered through Lisp systems.
- With Michael Reinfrank, Johan de Kleer, and Matt Ginsberg, *Nonmonotonic Reasoning. Proceedings of the Second International Workshop,* (Springer Verlag, Lecture Notes in Artificial Intelligence, number 346, 1988).
- With Christer Bäckström, *Current Trends in AI Planning,* (IOS Press, 1994).

Journals

In 1997 I started the *Electronic Transactions on Artificial Intelligence* and became its General Editor. This journal pioneered the concept of open reviewing using a scheme of two-stage peer review, consisting of an open (non-confidential) reviewing stage, followed by a confidential refereeing stage. In addition it used the principle of true open access, ie. with neither subscriptions nor author fees. The use of open access was unusual but not unique at the time, but the open reviewing approach was entirely new, and in fact it is only in the last few years that there has been a considerable growth of interest in such alternative peer review methods.

In 2001 I was elected co-Editor-in-Chief (with Ray Perrault) for the Artificial Intelligence Journal, which is the prime international journal in our field. I held this position for one term, until 2006. This was an important and interesting assignment, although with one major disappointment: I had hoped that some new approaches to scientific journal publishing could be introduced in the AIJ, but unfortunately this could not be realized.

I have also been an Editorial Board member for a number of journals, both within Artificial Intelligence and in neighboring fields, including in particular:

```
Artificial Intelligence Communications
Computational Intelligence
Fundamenta Informaticae
Journal of Applied Non-classical Logics
Applied Artificial Intelligence
Decision Support Systems
Information Systems
```

Selected publications

These publications have Erik Sandewall as the single author unless otherwise is noted. As a general rule I have not wanted to be included as co-author of those articles by my students where my contribution had only been as advisor, so co-authored articles are those where I have cooperated with peers or where I made a substantial part of the work myself.

Books

1) Erik Sandewall: *Features and Fluents.* Oxford University Press, 1994.

Journal Articles

2) A Planning Problem Solver Based on Look-Ahead in Stochastic Game Trees. In: *Journal of the Association for Computing Machinery,* volume 16 (1969), pages 364-382.

3) Torgny Groth, Werner Schneider, Jean-Claude Vuille, Erik Sandewall: Computer Simulation of Ferrokinetic Models. In: *Computer Programs in Biomedicine,* volume 1 (1970), pages 90-104.

4) Formal methods in the design of question-answering systems. In: *Artificial Intelligence,* volume 2 (1971), pages 129-145.

5) A Proposed Solution to the FUNARG Problem. In: *The SIGSAM Bulletin,* volume 17 (1971), pages 29-42.

6) Conversion of predicate-calculus axioms, viewed as non-deterministic programs, to corresponding deterministic programs. In: *IEEE Transactions on Computers,* volume C-25 (1976), pages 342-346.

7) Lennart Beckman, Anders Haraldsson, Östen Oskarsson, Erik Sandewall: A partial evaluator, and its use as a programming tool. In: *Artificial Intelligence,* volume 7 (1976), pages 319-357.

8) Programming in the Interactive Environment: The LISP Experience. In: *Computing Surveys,* volume 10 (1978), pages 35-71.

9) A functional approach to nonmonotonic logic. In: *Computational Intelligence,* volume 1 (1985), pages 80-87.

10) Nonmonotonic Inference Rules for Multiple Inheritance with Exceptions. In: *Proceedings of the IEEE,* volume 74 (1986), pages 1345-1353.

11) Towards a Logic of Dynamic Frames. In: *International Journal of Expert Systems,* volume 3 (1990), pages 355-370.

12) Knowledge-based systems, Lisp, and Very High Level Implementation Languages. In: *The Knowledge Engineering Review,* volume 7 (1992), pages 147-155.

13) Magnus Morin, Simin Nadjm-Tehrani, Per Österling, Erik Sandewall: Real-Time Hierarchical Control. In: *IEEE Software,* volume 9 (1992), pages 51-57.

14) The Range of Applicability for some Nonmonotonic Logics for Strict Inertia. In: *Journal of Logic and Computation,* volume 4 (1994), pages 581-616.

15) Towards a World-Wide Database. In: *Computer Networks and ISDN Systems,* volume 28 (1996), pages 1513-1522.

16) Towards the Validation of High-Level Action Descriptions from their Low-Level Definitions. In: *Artificial Intelligence Communications,* volume 9 (1996), pages 214-224.

17) Publishing and Reviewing in the ETAI. In: *Electronic Transactions on Artificial Intelligence,* volume 1 (1997), pages 1-12.

18) Logic-based Modelling of Goal-Directed Behavior. In: *Electronic Transactions on Artificial Intelligence,* volume 1 (1997), pages 105-128.

19) Cognitive Robotics Logic and its Metatheory: Features and Fluents Revisited. In: *Electronic Transactions on Artificial Intelligence,* volume 2 (1998), pages 307-329.

20) An Early Use of Continuations and Partial Evaluation for Compiling Rules Written in First-Order Predicate Calculus. In: *Higher-Order and Symbolic Computation,* volume 12 (1999), pages 105-113.

21) Mark S. Frankel, Roger Elliott, Martin Blume, Jean-Manuel Bourgois, Bernt Hugenholtz, Mats G. Lindquist, Sally Morris, Erik Sandewall: Defining and Certifying Electronic Publication in Science. In: *Learned Publishing,* volume 13 (2000), pages 251-258.

22) On the Design of Software Individuals. In: *Electronic Transactions on Artificial Intelligence,* volume 5 (2001), pages B:143-160.

23) Systems: Opening up the process. In: *Nature,* DOI 10.1038/nature04994 (2006),

24) A Review of the Handbook of Knowledge Representation. In: *Artificial Intelligence,* volume 172 (2008), pages 1965-1966.

25) Extending the concept of publication: Factbases and knowledgebases. In: *Learned Publishing,* volume 21 (2008), pages 123-131.

26) Exercising Moral Copyright for Evolving Publications. In: *ScieCom Info,* volume 6 (2010), pages 1-4 (*).

27) Defeasible inheritance with doubt index and its axiomatic characterization. In: *Artificial Intelligence,* volume 174 (2010), pages 1431-1459.

28) From systems to logic in the early development of nonmonotonic reasoning. In: *Artificial Intelligence,* volume 175 (2011), pages 416-427.

29) Maintaining Live Discussion in Two-Stage Open Peer Review. In: *Frontiers in Computational Neuroscience,* volume 6 (2012), pages 1-11 (*).

30) A Perspective on the Early History of Artificial Intelligence in Europe. In: *Artificial Intelligence Communications,* volume 27 (2014), pages 81-86.

(*) The asterisk marks journals where each article has page numbers from 1 and up.

Conference Articles

31) Lisp A, A Lisp-like System for Incremental Computing. In: *Spring Joint Computer Conference,* 1968, pages 375-384.

32) A Set-oriented Property-Structre Representation for Binary Relations. In: *Machine Intelligence 5,* 1970, pages 237-252.

33) PCDB, a Programming Tool for Management of a Predicate Calculus Oriented Data Base. In: *International Joint Conference on Artificial Intelligence,* 1971, pages 159-166.

34) Representing Natural-Language Information in Predicate Calculus. In: *Machine Intelligence 6,* 1971, pages 255-280.

35) An Approach to the Frame Problem, and its Implementation. In: *Machine Intelligence 7,* 1972, pages 195-204.

36) Deductive Search in a Semantic Net. In: *Symposium 'Organismic Information Processing',* 1973, pages 512-525.

37) Conversion of predicate-calculus axioms, viewed as non-deterministic programs, to corresponding deterministic programs. In: *International Joint Conference on Artificial Intelligence,* 1973, pages 230-234.

38) Ideas about Management of LISP Databases. In: *International Joint Conference on Artificial Intelligence,* 1975, pages 585-592.

39) Some Observations about Conceptual Programming. In: *Machine Intelligence 8,* 1977, pages 223-265.

40) A Description Language and Pilot-System Executive for Information-Transport Systems. In: *Conference on Very Large Data Bases,* 1979, pages 101-110.

41) Biological Software. In: *International Joint Conference on Artificial Intelligence,* 1979, pages 744-747.

42) Erik Sandewall, Claes Strömberg, Henrik Sörensen: A System for Communicating Residential Environments. In: *Lisp Conference,* 1980, pages 82-89.

43) Erik Sandewall, Göran Hektor, Anders Ström, Claes Strömberg, Ola Strömfors, Henrik Sörensen, Jaak Urmi: Provisions for Flexibility in the Linköping Office Information System (LOIS). In: *National Computer Conference,* 1980, pages 569-578.

44) Erik Sandewall, Claes Strömberg, Henrik Sörensen: Software Architecture Based on Communicating Residential Environments. In: *International Conference on Software Engineering,* 1981, pages 144-152.

45) Unified Dialogue Management in the Carousel System. In: *Office Information Systems Workshop,* 1982, pages 175-197.

46) Erik Sandewall, Sture Hägglund, Christian Gustafsson, Lennart Jonesjö, Ola Strömfors: Stepwise Structuring - A Style of Life for Flexible Software. In: *National Computer Conference,* 1983, pages 65-72.

47) Erik Sandewall, Ralph Rönnquist: A Representation of Action Structures. In: *[U.S.] National Conference on Artificial Intelligence,* 1986, pages 89-97.

48) Specification Environments for Information Management Systems. In: *International Federation of Information Processing,* 1986, pages 313-314.

49) The Pipelining Transformation on Manufacturing Cells with Robots. In: *International Joint Conference on Artificial Intelligence,* 1987, pages 1055-1062.

50) Future developments in Artificial Intelligence. In: *European Conference on Artificial Intelligence,* 1988, pages 707-715.

51) An Approach to Non-Monotonic Entailment. In: *Methodologies for Intelligent Systems, III,* 1988, pages 391-397.

52) Filter Preferential Entailment for the Logic of Action in Almost Continuous Worlds. In: *International Joint Conference on Artificial Intelligence,* 1989, pages 894-899.

53) Combining Logic and Differential Equations for Describing Real-World Systems. In: *Conference on Principles of Knowledge Representation and Reasoning,* 1989, pages 412-420.

54) Erik Sandewall, Giovanni Adorni, Hans-Hellmut Nagel, Monique Thonnat: PROMETHEUS Session. In: *European Conference on Artificial Intelligence,* 1990, pages 779-780.

55) Reasoning about the World as Perceived by an Agent. In: *European Conference on Artificial Intelligence,* 1990, pages 579-584.

56) Lisp as a Very High Level Implementation Language. In: *European Conference on Practical Applications of Lisp (EuroPAL),* 1990, pages 29-34.

57) The Role of Temporal Reasoning Subsystems in the Architecture of Autonomous Agents. In: *Artificial Intelligence in Engineering,* 1993, pages 3-6.

58) Systematic Assessment of Temporal Reasoning Methods for Use in Autonomous Agents. In: *Methodologies for Intelligent Systems, VII,* 1993, pages 558-570.

59) The Range of Applicability of Nonmonotonic Logics for the Inertia Problem. In: *International Joint Conference on Artificial Intelligence,* 1993, pages 738-743.

60) Getting Robots to Know Exactly What They are Doing. In: *European Conference on Artificial Intelligence,* 1996, pages 707.

61) Assessment of ramification methods that use static domain constraints. In: *Conference on Principles of Knowledge Representation and Reasoning,* 1996, pages 99-110.

62) Transition cascade semantics and first results for ramification. Preliminary report. In: *Conference on Formal Approaches to Practical Reasoning,* 1996,

63) Relating High-Level and Low-Level Action Descriptions in a Logic of Actions and Change. In: *International Workshop on Hybrid and Real-Time Systems,* 1997, pages 3-17.

64) The Logical Characterization of Goal-Directed Behavior in the Presence of Exogenous Events. In: *Logics in Artificial Intelligence,* 1998, pages 382-384.

65) Logic-Based Modelling of Goal-Directed Behavior. In: *Conference on Principles of Knowledge Representation and Reasoning,* 1998, pages 304-315.

66) Patrick Doherty, Gösta Granlund, Krzystof Kuchcinski, Erik Sandewall, Klas Nordberg, Erik Skarman, Johan Wiklund: The WITAS Unmanned Aerial Vehicle Project. In: *European Conference on Artificial Intelligence,* 2000, pages 747-755.

67) Use of Cognitive Robotics Logic in a Double Helix Architecture for Autonomous Systems. In: *Advances in Plan-Based Control of Robotic Agents,* 2001, pages 226-248.

68) High-level Design of Web Servers in Allegro Common Lisp. In: *International Conference on LISP,* 2003,

69) Erik Sandewall, Patrick Doherty, Oliver Lemon, Stanley Peters: Real-time Dialogues with the WITAS Unmanned Aerial Vehicle. In: *Annual German Conference on AI,* 2003, pages 52-63.

70) A Software Architecture for AI Systems Based on Self-Modifying Software Individuals. In: *International Conference on LISP,* 2003,

71) Reification of Action Instances in the Leonardo Calculus. In: *IJCAI Workshop on Nonmonotonic Reasoning, Action and Change,* 2007,

72) Artificial Intelligence Needs Open-Access Knowledgebase Contents. In: *[U.S.] National Conference on Artificial Intelligence,* 2008, pages 1602-1605.

73) Intelligent Software Individuals Based on the Leonardo System. In: *Symposium on Advances in Cognitive Systems,* 2011, pages 273-279.

74) The Leonardo System and Software Individuals. In: *European Conference on Lisp,* 2013, pages 18-24.

Articles in Collection Volumes

75) Formal Specification and Implementation of Operations in Information Management Systems. In: *Colloquium Programmeeromgevingen, MC Syllabus,* 1983, pages 125-148.

76) The Semantics of Non-Monotonic Entailment Defined Using Partial Interpretations. In: *Non-Monotonic Reasoning,* 1988, pages 27-41.

77) Erik Sandewall, Yoav Shoham: Nonmonotonic Temporal Reasoning. In: *Handbook of Logic in Artificial Intelligence and Logic Programming,* 1995, pages 439-498.

78) Reasoning about Actions and Change with Ramification. In: *Computer Science Today,* 1995, pages 486-504.

79) Underlying semantics for actions and change with ramification. In: *Spatial and Temporal Reasoning,* 1997, pages 146-154.

80) Coordination of Actions in an Autonomous Robotic System. In: *Reasoning, Action and Interaction in AI Theories and Systems,* 2006, pages 177-191.

81) The Leordo Computation System. In: *From Semantics to Computer Science,* 2008, pages 309-336.