Good Old Fashioned Automata Theory and the Agents World

Leona F. Fass

Carmel-by-the-Sea
P.O. Box 2914
Carmel CA 93921
<lff4[AT]cornell[DOT]edu>

Abstract

Our architectural theory is automata theory and our processes combine knowledge representation and automata-based reasoning, producing the "best possible" agent configurations that may realize specified behavioral goals. We establish that if a goal can be realized by a configuration of Web-dwelling agents then a distinguished Manager Agent, with suitable Local Closed World knowledge of agents' properties, may apply automata-based reasoning algorithms, selecting and configuring agents into a realizing device. The Manager also may reconfigure, adapt and optimize devices, relative to available resources. Real world and logical constraints may mean that a "best possible" result is an approximation; this is confirmed as acceptable by other researchers using a variety of reasoning processes we describe. We generalize our results to local groups of agents, not just those Web-dwelling, and describe both theoretical ("Personal Travel Assistant") examples and applications to practice.

Introduction

The world of agents has much in common with the world of classical, "good old fashioned," automata theory. Both areas involve synthesis of "devices" to meet behavioral specifications, fulfilling particular functions or achieving particular goals. We have found that automata-theoretic principles and procedures involved in abstract behavioral analyses and abstract machine syntheses provide a foundation for an agent configuration theory. That is, the processes involved in determining abstract devices to realize specified behaviors may be adapted to determining configurations of Web-dwelling agents that fulfill real behavioral goals. If a goal can be achieved on the Web, appropriate agents can be collected and configured to achieve it. By determining the relationship between a realizing device structure and its potential behavior, the "best possible" agent configuration may be found.

Thus our architectural theory is automata theory. Our processes involve knowledge representation and reasoning: knowledge about agents and their capabilities is represented for a supervisory, distinguished Manager Agent. That Manager effects automata-theory-based reasoning algorithms to select appropriate agent components and configure them into each specified behavior-realizing device. We have theoretical results relating to the discovery and construction of Web-dwelling agent configurations that fulfill specified tasks or realize behavioral goals. We describe such results for constructing "optimal" agent configurations, and an application of the theory to practice (in the traditional example of a Personal Travel Assistant). Other researchers have developed theoretical and practical results related to configuring problem-solving groups of agents; we discuss how our work relates to theirs. Research has produced many successful applications of agents to problems in practice. However, theoretical, logical and real world practical considerations can prevent problems from being solved by Web-dwelling agents. We discuss such cases as well.

Good (Automata) Theory Lasts Forever

We entered the field of AI with a background in theoretical computer science, where we had focused on formal problems of computational learning. Our research emphasized behavioral analyses, synthesis and inference processes for determining behavioral models from behavioral examples or samples. This work was grounded by classical switching and automata theory. I.e., by defining relationships between the components of a device's specified potential behavior and the device's necessary structural components, early theoretical research (Moore 1956, Myhill 1957, Nerode 1958) established existence of solutions to abstract machine problems and effective techniques to determine them. These included techniques for analyzing device and behavioral structure; synthesizing devices to produce or realize specified behaviors; and minimization processes to eliminate extraneous components, resulting in reduced "optimal" device forms. Results established that an appropriate finite behavioral sample was sufficient to effectively determine a finite, minimal behavioral model. Thus a finite device producing an entire specified behavior could be found from a suitable finite sample. No claim was made that the classical techniques of almost fifty years ago would be efficient (back then, today's concept of efficient computation was barely being formed). But once they established that a problem could be solved effectively, subsequent methods could be devised to produce results while conserving computational resources. The classical theory still provides a foundation for today's abstract device design problems and also is applicable to real machines. It has enabled us to devise theoretical inference techniques (for learning from samples or examples) that have been adapted to practical processes and applied to modeling and learning problems of AI.

We learned much about the Agents World within a theoretical context and began to see the relationships with our inference work. While participating in a symposium on logic-based program synthesis we attended presentations about Web-dwelling agents' problem-solving potential. Sheila McIlraith (McIlraith 2002) spoke of languages and processes developed for enabling agents to compose Web services automatically. Then, upon receiving a complex goal-directed request, the agents could select actions and transitions needed to fulfill the goal's component tasks (e.g., completing all phases in the planning and booking of a research trip). Richard Waldinger (Waldinger 2002) lectured on logic-based techniques to locate agents existing on the Web and configure them into temporary problem-solving systems. He and his colleagues have shown (Waldinger *et al* 2004) that available agents' capabilities could be described as axioms. When the axioms were provided to a theorem-prover, a theorem could be produced, effectively "gluing" appropriate agents together to achieve a specified behavioral goal (e.g., scheduling meetings in the face of geographic and temporal constraints). Alternately, if such a goal-oriented theorem could not be found, the process might reveal information inconsistencies and anomalies. We learned more about agents' problem-solving collaborations as realized by agent coalitions (Soh and Tsatsoulis 2002, Soh and Anderson 2004). There, traditionally, agents themselves determine the efficacy of working together to fulfill specific tasks. They may retain information about each other, including histories of joint successful endeavors. Then they may negotiate, argue or otherwise convince each other to form groups that achieve particular goals. Typically, a sub-optimal coalition is configured initially. Then it is adapted iteratively to become a relatively optimal device.

Our automata-theoretic grounding enabled us to view these examples as within the broader area of behavioral modeling. It became clear to us that the agents McIlraith described were involved in behavioral analysis to decompose a specific goal task into subtasks. They were involved in device synthesis as they interacted to fulfill the subtasks and achieve their behavioral goal. Device synthesis also applied to Waldinger's problem-solving agent constructions. Furthermore, the determination of anomalies, instead of behavior-realizing constructions, reminded us of procedures to test abstract or real devices for equivalence with their behavioral specifications. (E.g., the model checking processes that use finite-state "machines" to represent software behavior were first devised to verify behavioral correctness, but for most software properties they are effective for detecting anomalies and errors.) The agent coalition examples, while typically configured by self-directed agent components, reminded us of classical automata minimization techniques. There, beginning with a given device design, unnecessary components are eliminated iteratively in order to produce an optimal result.

The automata and switching results that formed the basis for our initial inference research involved abstraction and theory. The theory endured and has been applicable in numerous instances involving a multitude of interpretations. For example, the notions of "device" and "behavior," respectively, could be: an elevator control panel and the elevator's up-and-down trips; a finite sequential machine and its input *vs.* output; a grammar and the language it derives. Given the insights we have gained from agents research just described, we propose that the very same theory that relates behaviors and structures is applicable to configuring groups of Web-dwelling agents to fulfill specified tasks.

Relating Agents to Automata and Getting Results

We propose that automata theoretic concepts and processes will adapt to the Agents World and the tasks agents individually and collectively may fulfill. Theory makes it feasible to synthesize agent "devices" that "behave" as specified. The agents existing on the Web may be selected and configured into teams, groups or multi-agent systems, as appropriate, with resultant agent configurations constructed to realize specified behavioral goals.

In our earlier work on agent-related theory (Fass 2004-2006) we pointed out some of the correspondences we found between the Automata World and the Agents World. In the Automata World, devices are configured from states and the transitions among them. A transition to a specific state will occur as the result of some preceding behavior. A transition to a final state means that a behavioral goal has been achieved. In our view of the Agents World individual agents correspond to "states" and agent interactions correspond to "state-to-state transitions." A configuration of interacting agents corresponds to a device. Their fulfillment of a specified task corresponds to "reaching a final state." Configuring a collection of agents to fulfill a task corresponds to the automata theory concept of realizing a behavior. But how would this be done?

Given all of the agents in existence on the Web at any instant, each would be able to produce a particular behavior or fulfill a particular task (this could include a "null task"). Knowing the capabilities of each of the Web-dwelling agents, it would be possible to determine those that might be utilized to fulfill component subtasks of complex specified behaviors. For any complex task specified it would be possible (in theory) to find equivalence or "indistinguishablity" classes of all of these agents, relative to that specific behavioral goal. The equivalence relation could be "produce the same behavioral component of the specified task." (For example, a goal task might involve planning a trip from Carmel California to Palo Alto for a meeting at Stanford. A behavioral component might be the subtask of planning the initial part of the trip, from Carmel to Monterey. There may be many ways to do this: planning a car trip, planning a local bus trip, etc., which could fall within the same equivalence class of "getting to Monterey.") Of course for any task T there would be agents, perhaps most of them, that could have no role in fulfilling T or any of its subtasks. Relative to T these would all become part of a single "non-use" class when the indistinguishablity equivalence classes of agents are found (those familiar with automata theory would recognize that such an agent class would be similar to an automaton "dead state"). With knowledge of all of the agents capabilities and the task-related classes to which they each belong, it would be possible to synthesize a "device" out of agents fulfilling the constituent subtasks of the complex specified behavioral goal. Selecting an agent from each of the relevant classes and configuring them suitably into an agent system would be the construction technique. Only one representative of each agent class need be selected. The suitable configuration (e.g., agent interaction paths) of the realizing agent "device" would correspond to the configuration of subtasks within the complex specified behavioral goal.

With a choice of available agents to be selected from a known behavioral class, some criteria might be employed to determine which to choose. E.g., the agent that behaved "best" or most efficiently in the past might be selected. However, if that agent were engaged in fulfilling some other task and were not available, an alternate might be selected as the "best possible" behaviorally-equivalent component for agent "device" *available at the time*. Now, with knowledge of everything available on the Web it would be possible to upgrade a task-fulfilling configuration as new and "better" potential component agents were introduced. In any case, it would be possible to determine if an agent component of a potential realizing device proved unnecessary or useless, so that iterative adaptations could be effected until a result were "optimized," e.g., through minimization. Furthermore, if a group of agents within a configuration proved to work well together, it could be possible to remember which agents those were so that they could be reused in future.

This seems quite a bit to keep track of: agent capabilities; agent availabilities; agents' comparative skills; constructions of interaction channels among agents; behavioral goals; decompositions of tasks into subtasks, and so on. But we are developing *theory*. In (Waldinger *et al* 2004) the agent configuration process was effected by an overseeing theorem-prover that had been supplied with axioms. In the (McIlraith 2002) interpretation compositions were effected by software with logical decision-making capabilities, able to process special mark-up languages and ontology representations. In (Soh and Tsatsoulis 2002, Soh and Anderson 2004) most coalition configurations were determined by agents with their own informational databases, using case-based reasoning and reinforcement based on their interactions past. In our theory, we prefer agent selections and configurations determined by a distinguished Manager Agent (Fass 2005). The Manager may have knowledge of the other agents and be able to reason about them, select them and configure them into each specific task-fulfilling device. But the Manager will never try to reason about itself. Aside from the logical paradox that might arise, we agree with (Riley and Veloso 2004) that there is an advantage to separating an agent "coach" from the agent "team" it may advise. A distinguished Manager can work with different agent groups, using different configuration strategies; this wouldn't be feasible if the Manager joined a configuration itself and couldn't retain an objective overview.

On the Web new information is constantly introduced and existing entities often disappear. Thus no person or agent can ever be sure, when searching for Web-dwelling entities, that searching has been done "enough" and the best result has been found. To alleviate this problem, we may relax our demands on the Manager to keep track of "all of the agents in existence on the Web at any instant," even in a theoretical study. We will employ the (Heflin and Munoz-Avila 2002) concept of a Local Closed World (LCW) subset of the Web and let the Manager have LCW-knowledge. Then, within the LCW, the Manager can keep track of agent capabilities, availabilities, changing resources, etc. Within the LCW the Manager will know if a search has been "enough," and whether an entity sought within the LCW can "ever" be found. With these assumptions we can establish:

For every behavioral goal realizable on the Web (using resources available within a LCW), a distinguished Manager Agent with LCW-knowledge may effect the following techniques and achieve the following results:

(i) discovery of an optimal agent group (relative to the LCW) for fulfilling realizable specified tasks and configuring them into a device realizing the behavioral goal;
(ii) removal of agents extraneous or useless for a problem-solving configuration, retaining its behavior;
(iii) merger of agents behaving indistinguishably relative to a behavior specified;
(iv) "recycling" of successfully performing agent groups or subgroups, to be used in future behavior-realizing configurations when this proves to conserve resources (e.g., search time and space);
(v) discovery of minimized communication pathways among agents and agent systems that may realize a behavioral goal within the LCW;
(vi) expansion of a task-fulfilling agent configuration as a specified behavioral goal grows (and remains realizable within the LCW);
(vii) adaptation of a configuration as a behavioral goal, or constraints, may change (if still defining a behavior realizable within the LCW) ;
(viii) testing of an agent configuration (relative to its LCW) to see if it does produce a specified behavior.

To achieve the above, the Manager Agent would have the logical, analytic skills to decompose specified tasks into component subtasks; would be able to make comparisons and apply effective decision-making abilities; would have a dynamic knowledge base reflecting the capabilities and availabilities of the other agents within its (local) world. Matching task components with capable agents to fulfill such subtasks, a goal-directed agent group could be configured. Our theoretical approach attributes much intelligence and ability to the Manager Agent, but such are the qualities already exhibited by the proposed and working systems described by (McIlraith 2002) and (Waldinger 2002, Waldinger *et al* 2004) that originally inspired us. In our interpretation a distinguished Manager Agent may obtain the results above, using techniques of good old fashioned automata theory (adapted as described) by relating agent device structure to intended specified behavior.

However, there is a significant distinction between behavioral analysis and synthesis when applied to the Web, and when applied in automata theory. Automata theory is generally concerned with finding finite models of specified behaviors that may be infinite (e.g., a sequential machine that recognizes "all strings with at least 2 ***a***s and any number of ***b***s"). If there is a finite realization of a behavior, its components will be found. In automata theory if a finite realizing device is found, a minimal realizing device may also be found and may be considered optimal. But when considering the behavior produced by agents on the Web, there will always be just a finite set of agents to consider. Any behavior realized by a "device" configured from Web-dwelling agents will either be finitely-realized, or won't be realized by these agents at all. A behavior need not even be complex to be unrealizable. Perhaps no agent has yet been designed to fulfill one of its subtasks; perhaps the necessary component agents exist, but just aren't available. Furthermore, while there may be a concept of minimality, since the agents are real and produce real (not theoretical) behaviors, there may be some other accepted criterion for an optimal result. Thus we may also conclude:

If a specified behavioral goal can be realized by a configuration of Web-dwelling agents,

> then there will be an optimal (maybe minimal, maybe most time-wise efficient, or other defined "optimal" relative to the LCW) finite configuration of agents to realize that goal. A Manager Agent may construct the optimal result by configuring known agents, or by effectively adapting a sub-optimal configuration that is already known.

For example, if the Manager recognized that a component group of agents needed to fulfill a specific subtask had been used in the past (e.g., finding conversions of USD to GBP, or charging purchases on credit cards) the configuration could be recalled from its knowledge base [similar to the coalition agent reuse described in (Soh and Tsatsoulis 2002, Soh and Anderson 2004)]. If, say, our version of a Manager determined that a "better" agent to fulfill a subtask became available, e.g., one that worked faster, a substitution could be made.

We can illustrate some of the concepts described above with additional (traditional) travel planning examples. We may have a theoretical Personal Travel Assistant (PTA) that composes Web-dwelling entities into the needed components for planning our trips. The condition of "optimality" might include minimal travel time and/or minimal travel expenses. Individual agents will have representations of travel modes, timetables, pricelists, distance measures, calendars, schedules of meetings we need to attend, record of our personal preferences, etc. The agents may be "optimally" configured to produce the plan of: taking a low-cost local bus from Carmel to Monterey; taking another, nonstop, from Monterey to San Jose (SJ); then taking a Santa Clara Valley bus to Palo Alto. At certain departure times the optimal plan might involve relatively inexpensive "interlocking" airport shuttles (Carmel to SJ Airport; SJ Airport to Palo Alto). At other times the optimal plan would be the only possible plan: using an expensive taxi for the trip (at times/dates when public transportation isn't available). If, additionally, Carmel's only 2 taxis were being repaired, the behavioral goal might not be realized at all. If a behavioral goal changed, e.g., to go from Carmel to the SJ Convention Center for meetings, the PTA could "recycle" the successful Carmel-Monterey-SJ components from the previous trip plan. If the goal changed to travel for a meeting in NY the PTA might "recycle" its Carmel-SJ Airport plans and expand the configuration to include airplane trips from SJ Airport into JFK, etc. If an existing plan could be improved, an intelligent Manager could always make improvements within its resource scope.

Constraints and Successes

If a goal *can* be achieved on the Web, appropriate agents can be collected and configured to achieve it. By determining the relationship between a realizing device structure and its potential behavior, the "best possible" agent configuration may be found. However, there are many logical and real world constraints that may provide obstacles to perfect agent configuration and goal realization. There will always be problems of information completeness, quality and security that can't be controlled by "theory." Constraining a configuring Manager to LCW information may be realistic with respect to termination of searches. But due to the dynamic nature of the Web and all electronic media, a necessary or "best" potential configuration component may actually exist (perhaps outside of the LCW) but may never be found. Thus we must accept that behavioral-realizing agent configurations, even if LCW optimal, may really just be approximations to an ideal. Completeness issues imply that the configuration result may not be a "universal" optimal, just the best that could be found. Information accuracy also constrains the utility of results. E.g., our theoretical PTA may access schedules for buses from Carmel to Monterey but that doesn't mean that such information actually is accurately represented on the Web, in any electronic medium or otherwise. False information, whether unintentionally or maliciously provided, exists on the Web and in all other media too. Even if accurate information exists, adequately reflecting real world conditions at one instant, in the dynamic world, just as on the dynamic Web, things change. (Like everyone else who participates in Stanford meetings, we have many "traveling researcher" examples illustrating real world deficiencies of an otherwise optimally designed PTA.) Thus *human experience* has taught us to get local "eyewitness" evidence; question facts or double-check; and make contingency plans using common sense. Agents, individually or in goal-oriented configurations, don't (yet) have such human-level capability. Furthermore, while we expect "all knowledge" to exist on the Web and be available to agents, even if possible this could rarely reflect real world situations in real time. Thus not everything that's "needed to know" will be available to a Manager determining problem-solving configurations or to the task-fulfilling agent components. Not every real behavioral goal can be realized by configurations of agents. Still, if a goal can be achieved, we've provided a good old fashioned automata theory foundation for its achievement.

Related Work and Conclusions

Groundwork for practical development of a diversity of agents and agent systems, along with a range of potential and realized applications, was presented in the insightful (Jennings, Sycara and Wooldridge 1998). We consider the work on agents and the Web described by (Hendler 2001) to be foundational in its own way, so we have framed our agent theory within the context of the Web. But any result we have can be restricted to a more localized universe of agents for which a Manager Agent would have LCW knowledge. Then specified local behavioral goals could be realized. Many theoretical and practical researchers have found such localized approximating agent configurations, i.e., teams, coalitions or multi-agent systems, to produce acceptable results fulfilling specific tasks. Iterative adaptations often follow, to improve initial results. E.g., (Soh and Tsatsoulis 2002) apply adaptation of sub-optimal coalition results to obtain improved results that have been utilized successfully in military missile-tracking problems.

Supervisory "coaching" or other oversight entities, similar to our Manager Agent, have been employed by various researchers to effect configurations. Many instances are reported in (Tumer and Stone 2002, Soh and Anderson 2004) and we have described some such research above. Particularly relevant to our view, we found (Sen and Kar 2002) use "teacher" agents to manage and share information with "student" agents assisting them in varying decision-making methods. The "coaching" agent of (Riley and Veloso 2004) uses Markov Decision Processes to observe its team's past performance and determine an optimal policy for its future. While developed for the "soccer venue" many other general, real applications are possible. In (Guestrin, Venkataraman and Koller 2002) a value function is computed for a multi-agent system, based on an agent's state and the actions of some of the system's other agents. The resultant value helps the agents decide on what future joint actions to take. Luc Steels' agents (Steels 2004) develop communication skills adaptively and "human-like," to interact with each other in systems. Agent "competitors" are modeled using Bayesian techniques in (Garrido, Brena and Sycara 2002), where it is acknowledged the modeling method is only optimal relative to the information that is available. This is what we noted, when describing application of our automata-based agent configuration theory to solving problems on the dynamic real Web or in the dynamic real world: one does the best one can with the resources available.

Groundbreaking work by Stan Rosenschein (Rosenschein 1985) related theory to robots (situated agents, of course). Influential work of (Rosenschein and Kaelbling 1996) used "situated automata theory" enabling agents they considered to react to the agents' environments. The agents we consider react to what their Manager decides; the Manager has access to a local environment and all it entails. Still, having devoted much of our research effort to theoretical problems of computational learning, we were gratified to find that some of our own earlier agents theory has been utilized in real goal-oriented problems; (Fass 2004) has been applied to robot agent configuration problems for the real world (Vig and Adams 2007). Although we have always stressed that good theory provides foundation for practice, we were pleasantly surprised to receive such validation that our own automata-theory based agent configuration work has had practical applications. Other relevant research examples applying theory to configuring agent systems, both abstract or real, are reviewed in (Fass 2005, 2006).

Our good old fashioned automata theory approach to behavioral realization and problem-solving works perfectly when determining abstract devices in constrained theoretical problem domains. When applied to more realistic processes, such as configuring behavior-realizing "devices" composed of interacting agents, the results aren't likely to be perfect. There are many real-world and logical constraints that delimit the determination of agent configurations that otherwise might solve specified problems, realize possible specified behaviors or fulfill possible specified tasks. But if a goal *can* be achieved by an interacting group, or system, of agents within some (LCW) environment, then our theory establishes that the appropriate configuration will, indeed, be found. It may also be deemed an optimal configuration, relative to available resources. For many problems our theory provides a basis for finding "best possible" results, and we consider that a realistic outcome. Many other researchers we have cited, using varying mathematical techniques to configure collections of agents, acknowledge that sub-optimal results and adaptations are acceptable. Approximation is not the exception; in real world applications it is the norm.

Our architectural theory is automata theory, and our processes have combined knowledge representation and automata-based reasoning algorithms to configure agent components for each behavior-realizing device. Often these results are approximations. But good theory provides the foundation for determining the best approximations, and in practical problems those may be the best possible results of all. By establishing that problem solutions exist and may be found, theory lays the groundwork for techniques applied in practice.

Acknowledgment

We thank the anonymous referees for their comments, improving this work.

References

Dignum, V., Corkill, D., Jonker, C. and F. Dignum, 2004, *Agent Organizations: Theory and Practice*, San Jose CA, July 2004, AAAI Press WS04-02.

Fass, L.F., 2004, "An Automata-Theoretic View of Agent Coalitions," in (Soh and Anderson 2004): 18-21.

Fass, L.F, 2005, "Some Agent Theory for the Semantic Web," in *ACM Software Engineering Notes*, Vol. 30, No.5 (September 2005), abstracted, p. 46, full paper accessible *via* the ACM Digital Library.

Fass, L.F., 2006, "A Logical Look at Agents' Problem-Solving on the Semantic Web," Short Paper *Dedicated to the Memory of E.S. (Stewart) Bainbridge*, presented at the *21st IEEE Symposium on Logic in Computer Science*, within the *4th Federated Logic Conference*, Seattle WA, August 2006. Abstract appears at http://www.easychair.org/FLoC-06/LICS-day226.html.

Garrido, L., Brena, R. and K. Sycara, 2002, "On Measuring the Usefulness of Modeling in a Competitive and Cooperative Environment," in (Tumer and Stone 2002): 48-54.

Guestrin, C., Venkataraman, S. and D. Koller, 2002,"Context Specific Multiagent Coordination and Planning with Factored MDPs," in (Tumer and Stone 2002): 17-24.

Heflin, J. and H. Munoz-Avila, 2002, "LCW-Based Agent Planning for the Semantic Web," in *Ontologies and the Semantic Web* (A. Pease, R. Fikes and J. Hendler, Editors), Edmonton AB, July 2002, AAAI Press, WS-02-11: 63-70.

Hendler, J., 2001, "Agents and the Semantic Web," *IEEE Intelligent Systems*, March/April 2001: 30-37.

Jennings, N.R., Sycara, K. and M. Wooldridge, 1998, "A Roadmap of Agent Research and Development," *Autonomous Agents and Multi-Agent Systems*, Vol. 1, No. 1 (1998), Kluwer: 7-38.

McIlraith, S, A., 2002, "Automated Composition of Web Services," Invited Talk, AAAI Symposium on Logic-Based Program Synthesis: State of the Art and Future Trends, Stanford University, March 2002.

Moore, E.F., 1956, "Gedanken-Experiments on Sequential Machines," in *Automata Studies* (C.E. Shannon and J. McCarthy, Editors), Annals of Mathematics Studies, No. 34 (1956), Princeton University Press: 129-153.

Myhill, J., 1957, "Finite Automata and the Representation of Events," Wright-Patterson AFB, 1957, *WADC Tech Report* 57-624.

Nerode, A., 1959, "Linear Automaton Transformations," *Proceedings of the AMS*, IX (1959), 541-544.

Riley, P. and M. Veloso, 2004, "An Agent Coaching Agents," in *Supervisory Control of Learning and Adaptive Systems* (M. Rosenstein and M. Ghavamzadeh, Editors), San Jose CA, July 2004, AAAI Press WS04-10: 61-62.

Rosenschein, S. J., 1985, "Formal Theories of Knowledge in AI and Robotics," *New Generation Computing*, Vol. 3, No. 4 (1985): 345-357.

Rosenschein, S. J. and L.P. Kaelbling, 1996, "A Situated View of Representation and Control," in *Computational Theories of Interaction* (S. J. Rosenschein and P. Agre, Editors), 1996, MIT Press: 515-540. [Reprinted from *Artificial Intelligence*, 73 (1995): 149-173.]

Sen, S. and P.P. Kar, 2002, "Sharing a Concept," in (Tumer and Stone 2002): 55-60**.**

Soh, L.-K. and J.E. Anderson (Editors), 2004, *Forming and Maintaining Coalitions & Teams in Adaptive Multiagent Systems*, San Jose CA, July 2004, AAAI Press WS04-06.

Soh, L.-K. and C. Tsatsoulis, 2002, "Learning to Form Negotiation Coalitions in a Multiagent System," in (Tumer and Stone 2002): 106-112.

Steels, L., 2004,"The Evolution of Communication Systems by Adaptive Agents," in *Adaptive Agents and Multi-Agent Systems*, LNAI 2636 (2004), Springer Verlag, Berlin : 125-140.

Tumer, K. and P. Stone (Editors), 2002, *Collaborative Learning Agents*, Stanford University, March 2002, AAAI Press, SS-02-02.

Vig, L. and J.A. Adams, 2007, "Coalition Formation: From Software Agents to Robots," *Journal of Intelligent and Robotic Systems*, Vol. 50, No. 1 (2007), Springer: 85-118**.**

Waldinger, R., 2002, "Deductive Chat Lines for Multiple Agents," Invited Symposium Talk, abstracted in *Logic-Based Program Synthesis: State of the Art and Future Trends*, (B. Fischer and D. Smith, Editors), Stanford University, March 2002, AAAI Press, SS-02-05: 5.

Waldinger, R., Appelt, D., Fry, J., Israel, D.J., Jarvis, P., Martin, D., Riehemann, S., Stickel, M.E., Tyson, M., Hobbs, J. and J.L. Dungan, 2004, "Deductive Question Answering from Multiple Resources," in *New Directions in Question Answering* (M.T. Maybury, Editor), 2004, AAAI Press: 253-262.