

Neither Here nor There: Inference Research Bridges the Gaps between Cognitive Science and AI

Leona F. Fass

P.O. Box 2914, Carmel CA 93921
lff4 [AT] cornell [DOT] edu

Abstract

Cognitive Science (CogSci) and AI are addressed from the perspective of inductive inference research, specifically as applied to language learning. Language so represents intelligence that results bridge gaps between the fields. We give examples of rigorous results intractable for AI machines and humans; AI results humans find satisfactory; and AI-Hard problems with “good enough” solutions, adaptively obtained. We conclude that lack of *the human experience* may preclude machines from human thinking, but CogSci can help AI produce human-acceptable results. Conversely, CogSci can benefit when researchers study human processes to improve AI machines. We see no competition: cooperation will advance both fields.

Introduction

We have looked at instances of human intelligence from multiple perspectives, giving us much to say about the interactions of Cognitive Science and AI. Our work has brought us into contact with both research communities, but we do not classify ourselves as belonging (strictly) to either community. We are “neither here nor there”.

Beginning a research program as a theoretical computer scientist, we determined some formal mathematical characterizations of knowledge structure. We developed rigorous techniques to produce “perfect” models of knowledge by inference, from sufficient samples that could be observed. As a result, our mathematical inference work became classified as Learning Theory. The precise techniques that we developed could be implemented by machines, bringing our research within the range of Machine Learning and so, AI. Next we sought to expand our research to problems of the real world and real humans’ intelligence. We soon realized we must relax our rigorous approach and obtain results that were feasible. We could apply our formal inference techniques to human knowledge acquisition, getting results that were less than perfect. We just had to accept that the approximate results obtained would be “good enough”. At the current stage of our research we believe that real-world problems generally do have “good enough”, rather than precise, solutions. We have changed our perspective to be more consistent with Cognitive Science than with formal mathematical theory or with AI.

Inductive inference was the focus of our initial research, and we soon applied this approach to the area of grammatical inference and language learning [e.g., (Fass 1996, 2004), with backpointers to earlier relevant work]. We didn’t know that problems we considered would later

be classified as AI-Complete or AI-Hard; we simply wanted to solve them. When we began our research we used logic, algebra and automata theory to describe structured language production and processing. This mathematical approach enabled us to represent formal context-free (CF) languages as perfectly “learnable”. Our research involved language and its representation, so we often interacted with linguists, philosophers, cognitive psychologists, and neuroscientists [e.g., those represented in (Powers *et al* 1991)]. We began to see language structure no longer mathematically, but in a more realistic and less precise way. We saw learning as an approximating, incremental process, and began thinking of adaptation as a critical component of intelligence.

In connection with (Cohen and Oates 2001) we wrote a brief paper we entitled “Language Learning Determines All Learning” for we realized that what we’d found, based on our linguistic knowledge acquisition work, was applicable to numerous AI and cognitive processes. We now consider questions [posed in (Lebiere *et al* 2005/6)] relating to Cognitive Science, AI and AI-Hard problems from a largely language-oriented perspective, for that is an application area with which we have had much experience. But we believe that language so much represents intelligence, that here we can draw general conclusions about Cognitive Science and AI. We have no doubt that precise, rigorous techniques can ground both AI and Cognitive Science approaches and problems. AI and Cognitive Science can contribute to each other, and what we have found from our inference research can help bridge the current gaps between the fields.

Is Cognitive Science Relevant to AI Problems?

We would say yes. All the way back to the origins of AI there has been a goal of having machines achieve human-level intelligence, but we aren’t there yet (McCarthy 2004). Some would say the point of AI was to get machines to “think”. A realistic current view might be that computers should save humans from doing too much thinking, by solving problems that would normally take up much human time. This is consistent with the first of Herb Simon’s stated goals for AI [quoted in (AITopics 2006)]: to augment human thinking. The second of his stated goals is for computers and AI to understand how humans think. That is where Cognitive Science has an important role to play. Cognitive Science should help AI in developing more human-like, human-compatible processes, rather than

processes that are strictly brute-force. Better human-computer interfaces would be an obvious result. In the MIT Intelligent Room Project (Coen 1999) noted that instead of making computer-interfaces for people, a goal was to make people-interfaces for computers. Integration of computers into everyday life was to be “seamless”. That was where the project encountered so many problems that are AI-Hard: vision, speech understanding and, more so, integration of hardware and software subsystems that were to achieve such goals.

We believe that the human mind can be a model for AI processes even though there is a paradox. On the one hand, machines cannot duplicate critical human thought processes (e.g., creativity or “having gestalts”) and must have every step of a process anticipated (one way or another). On the other hand, machines can deal with more information for immediate access at specific times and can work much faster than humans on numerous problems. Machines can be “better than humans” in that sense. We believe that many problems that became AI problems were those that humans could not solve in reasonable amounts of time, and this was one reason that the classes of problems considered by Cognitive Science and AI diverged.

We cannot foresee a time when computers will “think” in exactly the same way as humans. But for AI, we believe it is the end product of processes that matter, not necessarily the means by which they are achieved. E.g., in our linguistic knowledge acquisition research, it really didn’t matter to us if a human, a mainframe computer, or a situated embodied agent each had precisely the same internal representation of a language (Fass 2004). We just wanted communication among the entities to be effective and accurate, with all parties “speaking the same language” and being “understood”. We do believe there are tasks so difficult, and communication problems so extreme, that the end results as well as methods may be constrained. But if there is a problem that human and machine each *can* solve, we do not care about their respective solution methods. We do care that satisfactory results will be obtained. Cognitive Science should help AI machines to interact with humans in a more natural fashion and to produce more human-like end results.

Are “Good Enough” Solutions Valuable?

We absolutely and heartily agree that they are. A lesson learned in our precise, theoretical research was that perfect solutions are very nice to get, but they can rarely be found in problems of the unconstrained real world. In fact, perfect solutions may not exist in real domains. A “good enough” solution may be the best outcome obtainable. It may be sufficient to achieve behavioral goals, and it may be adaptable to produce a better solution in future.

For example, in our linguistic knowledge acquisition research, we found a “perfect” way to represent language describable by formal CF grammars, thus providing a foundation for effective language learning. But, if a

language *were* describable that way, experimental processes sufficient to *find* the target representation would be intractable (except in the most trivial cases). The perfect solution would exist and could be found effectively, but *realistically* there wouldn’t be enough time or space to find it. Furthermore, in almost every real-world case there could *only* be an approximating result, using the CF (or any) descriptive model. Just considering the changing lexicon (not even evolving syntax), only a subset of knowledge to-be-acquired could ever be observed. Any models based on observations thus would have to be approximate. But such “good enough” results can be adequate for representation and communication needs. Thus they would be of value.

We couldn’t ever use our proven inference techniques to learn all of a Natural Language (NL), but we were able to obtain an adequate approximation that was of value when we actually had to use it. Needing to learn Italian rapidly before going to present a lecture in Florence, we found an “approximating” grammar and vocabulary based on short-term observations. It served our purposes, and could be modified as new examples were observed. The “good enough” approximation enabling us to communicate was “perfect” for just a subset of knowledge-to-be-acquired. Cognitive scientists and AI practitioners know that although nobody has produced the perfect model that a human or AI system might use to learn a real language (and surely no one could do this tractably), humans using “good enough” approximations *do* learn immense amounts of language that precise machine systems simply can’t.

None of the research reported in (Cohen and Oates 2001) or (Cohen *et al* 2004) claims to have *the* solution to problems of language learning or to any other knowledge acquisition problem. But researchers can find solutions that are adequate. All of the work we see describes tentative solutions, partial solutions, and plans for improvements and future research. In theory one may be able to “stop the clock” and represent “all available knowledge”, but in reality the environment and knowledge humans deal with are dynamic and can always change (Fass 1996). Once a model is obtained, it can be immediately out-dated. But it is obviously valuable to represent “most” knowledge. Thus we share the engineering view described by Bourquet: “But I’m an engineer. Finding a model of the reality adequate for the task at hand but still [being] conscious that the model is imperfect is the rewarding side of my work. There are two risks: [being] too imprecise and [being] too complicated” (Bourquet *et al* 2004). We conclude that approximate, good enough solutions are not just valuable; they may be all that one can get. Then subsequent observations may lead to improved results, getting closer to the (most likely unreachable) behavioral ideal.

Are Multi-Level Heterogeneous Approaches Beneficial?

Yes, but.... We say yes, if the goal is to characterize and attempt to replicate human thinking processes. Human

learning combines many forms of reasoning and processing. We believe part of human intelligence is the ability to decompose complex problems into subproblems that can be solved heterogeneously and, often, flexibly. Many aspects of human-level intelligence are described by (McCarthy 2004). A specialized technique may be appropriate in a real-world domain, to be integrated into a larger complex problem solution or life situation. A real-world intelligent human would determine that a specialized technique should be applied and would effect the integration. Machines in the broadest sense (e.g., including multi-agent systems) can try to mimic the human decomposition process and work on subproblems heterogeneously. A single, focused method might be efficient for specialized and limited cases, but not sufficient for general intelligence and learning. However, if a goal is just for a machine to produce a result for a problem that a human might also solve, with no concern for how (or why) the result is obtained, a single, focused approach that works efficiently and cost-effectively may be the best technique.

Is Adaptiveness an Essential Component of Intelligence?

Absolutely! We can think of little else that is more indicative of intelligence. [We don't consider the mindless "adaptiveness" of a plant when it turns toward the sun. Nor do we consider the "no-memory" reinforcement-learning (RL) variety of "adaptiveness". RL systems are very mathematical and, we believe, a resultant limitation is that they are not very real or human-like. When reading the (Harmon and Harmon 1996) example of a trial-and-error RL goal of learning to ride a bicycle by avoiding repetitive "negative reinforcement", we recalled our own human bicycling experiences. To avoid repetitive and *painful* falls, we learned it would be best to stop trying to ride. An RL system does not consider a history of experiences while achieving a specific goal. It never gets a global knowledge "big picture" nor does it learn why a goal should be achieved. We deem these to be major RL limitations.] We consider adaptiveness that involves: an *awareness* of environment; an ability to make comparisons; and the capability of making judgments and decisions, perhaps cognizant of non-local conditions. It implies some ability to predict what might happen once changes are made (McCarthy 2004). We consider adaptiveness that may include prior knowledge and all of the above factors. It is an essential component of learning and thus, intelligence.

In our language-learning example, we found a learner could infer an acceptable knowledge model, based on information that was available or observed. But when new information arose, a learner could adapt the representation and a better, more complete result could be found. This approach to the optimal does show the modeler/learner can improve. It demonstrates intelligence since the modeler or learner is deciding to correct or modify what is known, thus getting closer to acquiring "all knowledge".

In the Neural Theory of Language (NTL) Project at UC Berkeley (Chang and Maia 2001) similarly developed an Analysis-Hypothesis-Reorganization paradigm for discovering domain regularities and forming representations. NTL linguistic constructs followed once concepts were represented. New language constructs could be *adapted* from existing ones, reflecting an increasingly sophisticated conceptual level. The authors noted this to be consistent with a child learning more language as s/he may age. Of course it also reflects intelligence. Adaptation and intelligence are apparent in the learning research described by (Solan *et al* 2004). There the learner observed a body of knowledge and used inherent statistical information and structural similarities to determine patterns *incrementally*. In the research reported by (Apoussidou and Boersma 2004) learners who used an adaptable, *gradual* learning algorithm succeeded, compared to those using other techniques, who did not. Without supervision, but by observing and *adapting*, a learner acquired a model of knowledge. MIT Media Lab research reported by (Roy 2004) described development of situated embodied systems, with *evolving* cognitive architecture to support knowledge grounding. In highly mathematical work (Arz-Goncarowski 2001) used abstract category theory to help generate representations over sensory motor perceptions. This provided a computational basis for *adapting* representations as perceptions changed. The representations could then be applied to a range of behavioral and cognitive mechanisms. In real-world situations, of course, environments and perceptions do change. Thus adaptation is essential for intelligent living beings. In the case of intelligent artificial beings, Luc Steels' superb research (Steels 2001, 2002, 2004) has considered emergent and *adaptive* communication systems for agents. He has emphasized that language evolves, and that communication is an adaptive behavior. This is true for intelligent humans and true for his proposed systems of autonomous distributed agents: software, or humanoid situated robot, varieties. A goal of his research (Steels 2004) has been for the agents to self-organize communication systems similar to NL. An indication of their AI-type intelligence would be their artificial language *evolution*.

Are the Most Efficient Solutions Problem-Specific?

Generally yes, but they may not be very useful. Our experience has been that for a general class of problems, it may be possible to find a relatively efficient technique that takes sufficient processing into account. But for specific cases it might be possible to find a more efficient solution, reducing or eliminating some of the processes/processing involved. For example, in a grammatical inference process we determined a technique that would always work to produce a syntactic model for *any* member of an entire language class. In all cases the inference process utilized information from a bounded sufficient sample and a

corresponding set of experiments. But for *specific* well defined subvarieties within the class better, more efficient techniques and solutions could be found. The general solution, while efficient, had to take worst cases into account. Specific-case solutions could be simplified considerably.

Now, if an intelligent being began with the complex general technique and “collapsed” it to find the more efficient specific-case procedure, much intelligence would be demonstrated and much knowledge could be gained. But if the intelligent being began with the specific case procedure, perhaps no problems other than the special one might be solved. Scaling up may be difficult, or might be impossible: a special case solution just might not generalize. (E.g., we could use our CF grammatical inference techniques to infer special-case regular language grammars, and we could simplify our general techniques to find a more efficient process for the regular case. But if we began with the special regular case technique it would not be applicable to arbitrary strictly CF cases. Modification to handle such cases would *not* be trivial.)

In relation to current AI-Hard problems, of course subproblems can be solved without the vast amount of information or processing needed to solve general cases. For example (Dale 2002) noted that while it is AI-Hard for machines to solve a problem involving human-level language understanding and communication, smaller parts of the problem, such as building functional telephonic voice-recognition systems, can be solved. (We, ourselves, have noted that such specific systems have improved significantly in recent years.) Rather than “chasing intelligence” he suggested that “smarter and shallower” should be the current aim. John McCarthy (McCarthy 2004) has opined that human-level NL understanding and processing are not yet achievable and the “little problems” solved won’t scale up. He believes AI *will* achieve human-level intelligence, but will get there when a fundamental, new way of thinking about these problems is applied.

We believe that if a provably correct general method can be determined, it establishes a foundation that grounds subsequently simplified specific methods. Then, finding more efficient specific solutions is something worth working towards. But if an intelligent being starts with an efficient specific solution, a satisfactory result might be obtained for just that case. At least there *will* be a result, but it may be all that the being will ever know.

Are Specialized, Modular Components a Reasonable Approach to General Intelligence?

Yes, and no. We believe that general intelligence includes the interaction, interconnection, and coordination of modular components. Specialized individual components solve subproblems and fulfill subtasks that arise as part of an intelligent being’s environment and existence. But an intelligent being has the capability to *coordinate* these component parts. This is our belief. It is consistent with the (Dale 2002, McCarthy 2004, and Steels 2004) remarks

relating human-like intelligence to the *entire* human experience. It is confirmed by the Intelligent Room Project’s discovery (Coen 1999) that integration of modules was the project’s most difficult work. Specialized modular systems such as vision, planning and speech all work together in humans, as part of human-level intelligence.

Based on our learning research, we believe that intelligent humans (not necessarily consciously) decompose problems into simpler ones to solve them. Subsolutions are then recomposed to produce solutions to the greater problems. This may involve biological factors, such as learning to put one foot in front of another to walk or developing hand-eye coordination. It may involve intellectual skills, such as learning to prove theorems. Or it may involve many aspects of the human experience, such as learning to talk, completing the steps needed to get to an office at Stanford or, even, baking cookies. A goal is analyzed and experience determines the tasks that must be, and are, fulfilled to achieve that goal. An intelligent being “remembers” the larger goal while fulfilling the modular tasks. Numerous sensory-motor and cognitive processes may be called into play.

When discussing “efficient problem-specific solutions” above, we remarked that they may not be very useful. We might say the same for an *uncoordinated* “modular” approach. It might be fine for AI to develop a “smart shallow” modular component that can do one thing very well, such as the systems described by (Dale 2002). At least that component can solve *some* problem and be of some use. It may be impressive for AI to develop a modular process making it appear as if a machine or agent has acquired a human-level skill. An example is the controversial recent (Cilibarsi and Vitanyi 2005) study that shows machines can “learn” NL semantics from analysis of Google pages. We find this work enlightening and useful, but bloggers have irately and quite rightly responded: “there is more to language than words!”

A modular process can produce nice, informative results. But if it does not scale up, and it is not part of a coordinated system of components (such as overseen by the human brain and mind), we don’t believe it is representative of general intelligence. We would say the same about human autistic savants who have astonishing skills in limited areas --- such as solving multiplication problems, or knowing and producing dates of historical events. Their demonstrated abilities do not represent general intelligence; they illustrate a unique modular skill.

Can Artificial Intelligence Contribute to Our Understanding of Human Cognition?

Yes, even though the fields have diverged. This divergence has occurred, in part, because of the difficulty of machines achieving human-level thinking and also, in part, because of advances in the technological aspects of AI. As we indicated in our discussion of Cognitive Science’s possible

contributions to AI, problems AI has been approaching often have been those that humans could not solve feasibly. In such cases, humans have been outdone by machines. However, we believe that a *human-like* version of intelligence may be used to design some AI systems. Then designers may analyze more about people's thinking. They may try to get machines to behave the same way or, at least, *to produce the same results* as humans. In this way, an AI-based analysis of human thinking will be a contribution to our understanding of cognitive humans.

For example, in our language learning work, we found the best way to represent language samples for inference was structurally. This representation (suggested to us by Leon Levy and Aravind Joshi) reflected the inherent "grouping structure" attributed to human-level language processing by many psycholinguists and neuroscientists. We then used the human-like representation to infer characterizing grammars in an automated way. We intended this to be a technique for machine learning of language, and it was very rigorous. Later we realized the technique could produce "good enough" human-learnable models of NL. As we have described, we used our own procedures to learn enough Italian to successfully negotiate our way around Florence (and environs). When we'd tried to learn sufficient Italian from a phrase book, or by listening to tapes, we did not learn enough to function or to understand what we had "learned". But when we used the psycholinguistically motivated AI technique, we thought more about the *human* way of learning language. Our "good enough" Cognitive Science result served us well. Virtually every language-related AI paper in (Cohen *et al* 2004) addresses some human cognitive process(es): hearing words, observing environments, understanding, making decisions, etc. Each may lead to advances in understanding the human mind.

The AI research related to the developing Semantic Web often analyzes human cognitive processes. A goal is for Web-navigating agents to take over human tasks and produce end results acceptable to human users. This has led to much useful research: gathering data; observing; and analyzing how humans act and think when they, themselves, use machines and the Web (Staab 2002, Steels 2002, Fass 2005, Cilibarsi and Vitanyi 2005). A positive outcome of rigorous Semantic Web development should be development of machines and systems that are more compatible with the cognitive processes of human users.

The artificial language evolution research described in (Steels 2004) is clearly AI, for he is working with autonomous agents. But to develop his processes he is examining the relationship between language conceptualization and environment; studying interactions between infants and caregivers leading to "learning how to mean"; and looking toward a Theory of Mind. This work is an impressive instance of developing AI processes for artificial beings by rigorously analyzing the cognitive processes of real human beings.

There is no question that the AI research just described can lead to many advances. Even if AI machines never reach human-level intelligence or produce human thinking precisely, we strongly believe that AI can contribute to understanding human cognition.

Conclusions

We have examined techniques and processes of AI and Cognitive Science from our background in theoretical inductive inference research and its applications. As a theoretician, we developed precise and rigorous language representation and linguistic knowledge acquisition techniques. They were applicable to Machine Learning problems and soon became classified as AI. But we had to relax our demand for "perfect" results as we sought to apply our approach to real-world problems. There we found we could obtain realistic "good enough" results, consistent with Cognitive Science. The problems of language representation and processing we considered, in the NL case, turned out to be AI-Hard. Our experience has given us a foundation for assessing others' research and opinions, and determining our own opinions about the future of both fields: Cognitive Science and AI.

We don't know that machines will ever achieve human-level intelligence, in the sense of thinking exactly as humans do. Even if machines could be embedded in the real world the way humans are, we don't expect machines to achieve a *precisely* human level. But we do believe there are problems AI will solve, producing human-like, human-acceptable results. We don't believe AI machines can ever replace humans, in all their cognitive processes. There is too much need for the *human experience* (e.g., seamless sensory integration, environmental awareness, social aspects, contextual considerations, predictions based on prior knowledge, etc.) to achieve cognition at a *human* level. We don't believe it will ever be possible to solve all human problems without sometimes consulting humans and this, we feel, is how it *should* be. There are many cases where AI can indicate to humans that something is worthy of a careful second look [as in "fair use" problems (Felten 2003), or tedious link analysis record-searching to detect possible crimes (Fass 1998)]. We don't want machines to finalize actions based on factors requiring human experience, and common sense, to assess.

Still, we do believe that Cognitive Science can help AI to develop systems that are more useful to humans, more compatible with their needs, and easier for them to interact with and understand. And we do believe that AI can help Cognitive Science by looking at human-level processes rigorously, in an attempt to duplicate their end results. E.g., there has been extraordinary improvement in the e-commerce voice-broker systems we recently have used. [Our current assessment confirms the earlier voice-recognition-system assessment by (Dale 2002), as we mentioned above.] Enough research by real NL scientists has been applied to this area to produce systems that

actually do work satisfactorily. Even we have the impression that the systems “understand” us. We used to shun such systems as “too stupid”. But with their current level of “expertise”, they are finally “bright enough” for us to use. We consider these improvements in language processing as indicative and illustrative of many other cognitive process advances: “language learning determines all learning”.

From our own “neither here nor there” research and our familiarity with the work of others, some of which we have reviewed above, we can see practical potential interactions of Cognitive Science and AI. We don’t view Cognitive Science and AI to be in competition. Rather, the areas can contribute to each other and their cooperation can advance both fields.

Acknowledgments

Symposium referees’ thought-provoking remarks led us to prepare this expanded final paper. Congleton Architect AIA and BMR Construction Management provided technical support for our current research.

In Tribute

We are grateful to the late Leon S. Levy for encouraging us to investigate grammatical inference problems, long ago, and for encouraging us to persevere.

References

- AITopics, 2006, “AI Topics: Cognitive Science” website, <http://www.aaai.org/AITopics/html/cogsci.html>, Jan. 12, 2006.
- Apoussidou, D. and P. Boersma, 2004, ”Comparing Different Optimality-Theoretic Learning Algorithms: The Case of Metrical Phonology” in (Cohen *et al* 2004): 1-7.
- Arzi-Gonczarowski, Z., 2001, “Generating Grounded Representations: A Mathematical Schema”, in (Cohen and Oates 2001): 4-9.
- Bourquet, J.M., Hartley, R. and E. Ozkural, 2004, posts Re: Zenkins paper on Cantor, Sept 30-Oct 2, 2004, <http://coding.derkeiler.com/Archive/General/comp.theory/2004-10/0017.html>
- Chang, N.C. and Maia, T.V., 2001, “Grounded Learning of Grammatical Constructions”, in (Cohen and Oates 2001): 10-15.
- Cilibarsi, R. and P.M.B. Vitanyi, 2005, “Automatic Meaning Discovery Using Google”, <http://arxiv.org/abs/cs.CL/0412098>, March 15, 2005.
- Coen, M.H., 1999, “The Future of Human-Computer Interaction, or, How I learned to stop worrying and love my Intelligent Room”, in *IEEE Intell Sys*, March/April 1999: 8-10.
- Cohen, P. and T. Oates (Eds), 2001, *Learning Grounded Representations*, Stanford Univ., March 2001, AAAI Press SS 01–05.
- Cohen, P., Clark, A., Hovy, E., Oates, T. and M. Whitbrock (Eds), 2004, *Language Learning: An Interdisciplinary Perspective*, Stanford Univ., March 2004, AAAI Press SS 04–05.
- Dale, R., 2002, “The Future of Smart Machines”, Macquarie Researcher, Sept 2002 Issue, <http://www.pr.mq.edu.au/researcher/showitem.asp?ItemID=152>
- Fass, L.F., 1996, “Minimal’ Language Learning”, in *Computational Issues in Learning Models of Dynamical Systems*, AAAI Spring Symp. Notes, Stanford Univ., March 1996: 68-71.
- Fass, L.F., 1998, “Inductive Inference and Link Analysis”, in *Artificial Intelligence & Link Analysis*, Orlando FL, Oct. 1998, AAAI Press, FS 98-01: 35-37.
- Fass, L.F., 2004, “Language, Mathematics and Other Dangerous Things”, *ACM SIGACT News*, Vol.35, No.1 (March 2004): 74-79.
- Fass, L.F., 2005, “Ubiquity and Common Sense for the Web”, in *Joint Conf. on Info. Sciences/Intl Symp. on Web Intelligence and Security*, Salt Lake City, Utah, July 2005: 1433-1436.
- Felten, E.W., 2003, “A Skeptical View of DRM and Fair Use”, *CACM* (Vol. 46, No. 4) April 2003: 56-59.
- Harmon, M.E. and Harmon, S.S., 1996, “Reinforcement Learning: a Tutorial”, <http://www.nbu.bg/cogs/events/2000/Readings/Petrov/rltutorial.pdf>
- Lebiere, C., Wray, R., Jha, K., and P. Weinstein, 2005/6, “Between a Rock and a Hard Place: Cognitive Science Principles Meet AI-Hard Problems”, CFP 2006 Spring Symp., <http://www.aaai.org/Symposia/Spring/ss06symposia.php#ss02>, Sept. 2005.
- McCarthy, J., 2004, “Roads to Human-Level AI?”, text of <http://wi-consortium.org/pdf/McCarthy-humanlevelAI.ps>, Sept. 2004.
- Powers, D., Harnad, S., Rayner, M. and L. Reeker, (Eds.), 1991, *Machine Learning of Natural Language & Ontology*, AAAI Spring Symp. Notes, Stanford, Univ, March 1991, released as German AI Institute, Kaiserslautern TR DFKI D-91-09.
- Roy, D., 2004, ”On Grounding the Content of Language” in (Cohen *et al* 2004): 57-61
- Solan, Z., Ruppin, E., Horn, D. and S. Edelman, 2004, ”Unsupervised Context Sensitive Language Acquisition from Large, Untagged Corpora” in (Cohen *et al* 2004): 61-65.
- Staab, S., 2002, (Ed.), “Emergent Semantics”, Trends & Controversies, *IEEE Intell Sys* (Vol. 17, No. 1), Jan/Feb 2002: 78-86.
- Steels, L., 2001, “The Role of Language in Learning Grounded Representations”, in (Cohen and Oates 2001): 80-83ff
- Steels, L., 2002, “Language Games for Emergent Semantics”, in (Staab 2002): 83-85.
- Steels, L., 2004, “The Evolution of Communication Systems by Adaptive Agents”, in *Adaptive Agents and Multi-Agent Systems* LNAI 2636. Springer Verlag, Berlin, 2004: 125-140.