

AAAI 1991 Spring Symposium Series Reports

■ *The American Association for Artificial Intelligence held its 1991 Spring Symposium Series on March 26–28 at Stanford University, Stanford, California. This article contains short summaries of the eight symposia that were conducted.*

Argumentation and Belief

Traditionally, research approaches to building models of argument knowledge and belief have differed in terms of objectives, scope, and methodology. For example, in rhetoric, the grounds-warrant-claim model has been used to analyze the structure of arguments. In psycholinguistics, researchers have analyzed the discourse structure of expository text by applying theories of discourse and schema coherency. In the field of logic, the emphasis has been on establishing axiomatic systems for deducing consistent beliefs. In AI, researchers have developed systems involving truth maintenance and evidential reasoning as well as models of legal reasoning, structural models of argument and discourse, and knowledge-based models of argument comprehension and generation.

A forum for multidisciplinary discussion of these approaches was provided last spring by the Symposium on Argumentation and Belief. The objective of the symposium was to further understanding of the issues, current status, and future directions of research in argumentation and beliefs. To facilitate the exchange of ideas, the presentation of each paper was followed by a commentator's response and a question-answering session. Papers were organized into five sessions that yielded insight into the following areas: representation of belief and argument knowledge, processes of argument comprehension and argument generation, the role of domain knowledge in argumentation, the role of memory in argumen-

tation, methods for assessing relative strengths and weaknesses of arguments, process of persuasion and belief change, the role of planning in argumentation, the role of case-based reasoning in argumentation, causal reasoning, analogical reasoning, and the teaching of argumentation skills. The symposium also included five panel discussions that dealt with issues involving the following areas: logical and philosophical models of argumentation and belief, methodology of research in argumentation and belief, representation of belief and argument knowledge, belief formation and learning through argumentation, and comprehension and generation of natural language in arguments.

The research described at the symposium demonstrates that designing systems capable of arguing requires characterizing interrelationships that exist between planning, reasoning, language comprehension, and language generation. For example, people reason about whether a goal should be achieved and use reasoning to justify beliefs about the efficacy, side-effects, and cost of selecting and executing plans. People also possess planning information on how to select and apply argument strategies. At the same time, human experts are able to not only present and justify their beliefs on possible courses of action but also understand opposing beliefs and argue persuasively for or against various positions. Given the complexity of these interrelationships, it is all the more important to continue developing and experimenting with whole, functioning prototypes, so that both the strengths and weaknesses of the proposed theories of argumentation and belief can be tested and revealed. The major benefit derived from the symposium was the multidisciplinary understanding of fundamental issues in knowledge representation, knowledge organization, and knowledge application that

must be addressed within any intelligent computer system capable of understanding and engaging in arguments.

Sergio J. Alvarado
University of California at Davis

Composite System Design

Presentations and discussions at this symposium focused on problems that were either real world and difficult or somewhat smaller, simpler, or well understood yet able to illustrate concepts and methods being proposed as common to composite systems and their design. To cope with the diversity of those attending, the workshop followed a hybrid strategy: pre-planned presentations and a survey of attendees to establish interests on which to spend the remaining time.

Composite Systems as the End Product of Design: Martin Feather (USC/Information Sciences Institute [USC/ISI]) introduced the concept of responsibility assignment within closed systems. The general notion is that for composite systems, it might be possible and advantageous to describe the properties desired of the system as a whole and use this point to begin to derive the properties required of the individual system components, which could then be implemented alone.

Stephen Fickas (University of Oregon) described his group's attempt to rationally reconstruct several existing composite systems using Feather's model. The preliminary results from these design exercises were that (1) a formal, transformational approach to responsibility assignment seems tractable in a tool that assists a human designer in generating designs and (2) formalizing and automating the diverse set of evaluation knowledge necessary to select a design is less feasible. Fickas argued that this last problem was

directly tied to composite system problems.

Eric Dubois (University of Namur) pursued the concept of responsibility, showing how it could be formalized in terms of deontic logic and incorporated within a requirement language.

Any nontrivial composite system or composite activity must involve a multitude of issues in addition to those directly linked to the properties of the system being designed. Les Gasser (University of Southern California [USC]) elaborated on these phenomena and outlined the structure of the HITOP Project that is studying the design of flexible manufacturing cells. This talk set the stage somewhat for what was to be a major theme at the workshop: designing to fit a tightly constrained context versus designing from a clean slate.

Eli Gerson (Tremont Institute) focused on the components of a composite system, outlining what it means to be an individual agent, whether one individual can be substituted for another, and how one can recognize the role an agent is playing in a composite system.

Fickas summarized his study of train accidents when viewed using Feather's model. Fickas noted that many components of modern train systems could be traced directly to overcoming one or more of a set of failure classes.

Chris Owens (University of Chicago) suggested that many complex systems have some default activity prescribed in cases of failure. He raised the questions of who detects and diagnoses failures, how they get communicated, and how they get repaired. Preestablished procedures seem to be compiled versions of some planning activity that considers these concerns.

The emerging field of coordination science also seems to relate to composite system design. Kevin Crowston (Massachusetts Institute of Technology [MIT]) outlined this relationship from a management perspective and stated that we should examine a variety of systems to determine commonly recurring coordination mechanisms as well as choice criteria that select from among them.

Distributed AI also has much in common with composite system design. Phil Cohen (SRI) presented formal notions of commitment, which permit the representation and reasoning about aspects of teamwork, coordinated action, dependability,

dialog, interaction, and so on. Similar concepts were the foundations of the work Yoav Shoham (Stanford University) presented on agent-oriented programming, wherein beliefs, commitments, and mental state are explicitly dealt with.

David Novick (Oregon Graduate Institute) outlined the use of speech act models to represent communicative acts such as the taking of turns to perform some activity, acknowledgments, and metalevel reasoning.

Finally, Carl Hewitt (MIT) posited the need for a structuring mechanism termed Orgs as a way to represent groups at all levels of abstraction.

Design as a Composite System Process: K.C. Burgess-Yakemovic (NCR) reported on experiments to apply the IBIS and GIBIS, approaches to recording design information. The essence of these experiments is the recording of issues, positions, and arguments that arise in discussions among designers in group meetings.

Keith Werkman (IBM) addressed the need to integrate the different evaluation criteria in civil engineering. He proposed a type of blackboard structure where the evolving design resided. Bill Robinson (Oregon) focused on the role of negotiation in multiperson design, particularly the need to represent and reason about the possibly conflicting goals of different system users or builders.

Rose Dieng (INRIA) considered knowledge acquisition from multiple specialists-experts from possibly different domains, including the end users of the system under development. She argued that any system that helps in this phase of activity is itself a composite system.

Application of Composite System Design to Real-world Tasks: A major question was whether we will ever have the luxury of designing a system from scratch, and if we don't, what this situation implies for composite system design. This discussion was referred to as "greenfield" versus "brownfield," an analogous comparison between building on a new site and building in an already developed location.

Ashok Goel (Georgia Tech) suggested that case-based reasoning can be integrated into a design tool, allowing one to reason about composite design problems in a complex domain even when one lacks a first-order domain theory. Goel discussed ARCHIE, an architectural design system using a case-based model.

Rob Helm (University of Oregon) posed an interesting challenge, suggesting that for any domain for which we have enough knowledge to do composite system design and critiquing, we must already have essentially completely codified all domain design options. Helm suggested that either we know enough about a domain to do (routine) design from tables, making a full composite system design tool exercise unnecessary, or we don't know enough about a domain to formalize it, and hence, (nonroutine) composite system design fails. This statement was considered an attack on AI and design, prompting heated discussion on realistic AI and design models and the use of brownfield design models.

Lewis Johnson (USC/ISI) summarized the need for manipulating a large amount of knowledge in designing composite systems and outlined ongoing work to capture and use such knowledge. He noted the need to reverse engineer composite system concepts, arguing that the current system should guide the retrospective rediscovery of design decisions.

Walt Scacchi (USC) reported on a large software development example with 40,000 requirements! As expected, many were conflicting, raising the question, Do real-world designers need sophisticated design tools, or do they simply need information management systems that store and track requirements, specifications, and designs, no matter how informal?

Sol Greenspan (GTE) used telephone services to illustrate a rapidly evolving and competitive domain that is clearly a composite system. He discussed the need to comply with standards, retain vendor independence, be able to use third-party service development, and maintain security and safety.

The number of hard problems raised at the workshop suggests that composite system design is too broad a research perspective. To satisfy everyone, we need an automated tool that can formally reason about responsibility, commitment, reliability, and motivation. Further, the tool should support nonroutine design in a brownfield domain where one must address social issues, organizational issues, laws and standards, and unavoidable failures. The tool itself must support multiperspective design and allow conflicting requirements and help sort them out.

Such a tool is far beyond our grasp.

However, if asked to choose one topic as a focus, the brownfield issue would be argued for: How do we model the existing infrastructure when attempting to add a new component or redesign an existing one? Can we reacquire or reverse engineer such a model if it is lacking? Can we evaluate what is changeable in the current system and at what cost? Can we analyze or simulate the effect of adding a new agent to a set of existing agents? If there is a second workshop, these questions should be addressed.

Stephen Fickas
Martin Feather
University of Oregon

Connectionist Natural Language Processing

For many years, both cognitive models and applications of natural language processing have used representations derived from symbolic logic from philosophy, linguistic theories of syntax, and AI. Recent results have led some researchers to propose that connectionism is an alternative to logic-linguistic-AI approaches to natural language processing. One goal of the symposium was to demonstrate, positively or negatively, that connectionism provides new insights into, or capabilities for, natural language processing.

The symposium was attended by 60 researchers from the United States, Germany, Japan, and the Netherlands. Work in connectionist natural language processing (CNLP) is done around the world, but the representation from various institutions was not balanced. Attendees from institutions on the West Coast of the United States accounted for 19 of the participants. The symposium also contained a balance of papers from graduate students and more established researchers, and the group maintained an atmosphere where the graduate students felt free to criticize the work of senior members of the community.

Natural language processing encompasses a range of phenomena, and one of the strengths of this symposium was the range of language problems, all of which were being attacked with connectionist techniques. From speech understanding to script application, all the researchers were exploring the usefulness of a common set of methods.

One feature that stood out among the papers was the lack of applications. Although applications were present, from call direction to information retrieval, most of the papers dealt with (1) cognitive science issues (for example, relating connectionist natural language to traditional linguistics, psychology, philosophy of mind) and (2) mechanisms for allowing networks to perform various operations, such as make dynamic inferences or encode recursive sentence structure.

In the relationship to linguistics, there was certainly no consensus. Most of the researchers at the symposium considered connectionism to be completely different from traditional linguistics, but there was some discussion that connectionism did fit with linguistic theories that are based on constraints as opposed to rules.

A recurring debate in many sessions was how much structure should be built into a network and how much should be learned. At the extreme end of the nonstructured side of the argument, a number of common techniques have emerged. Simple recurrent networks and recursive autoassociative memories are now considered standard tools for dealing with structured data in a relatively unstructured network. In the area of structured networks, no standard techniques have emerged. Various systems use everything from a phased clock to high-rank tensors for representing structured data.

Two trends did emerge. The first new trend is an increase in modularity, even among those systems based mostly on learning. This trend toward modularity came from both the cognitive and application-oriented research. However, increased modularity probably means more for applications because it reduces training time. The second trend was the move toward more *situated learning*, placing a network in an environment (either simulated or real) and letting it learn to perform a task that requires linguistic behavior. This trend will have an equal impact on applications and cognitive science. Applications can often benefit from learning systems that operate online, and situated learning makes a much better case for a cognitive model.

Charlie Dolan
Hughes Research Labs

Constraint-Based Reasoning

Constraint-based reasoning (CBR) is a paradigm that unifies many traditional areas in AI. It encourages the formulation of knowledge in terms of a set of constraints on some entities without specifying methods for satisfying such constraints. Many techniques for finding partial or complete solutions to constraint expressions have been developed and successfully applied to tasks such as design, diagnosis, truth maintenance, scheduling, spatiotemporal reasoning, and user interface design.

The symposium brought together a diverse group of researchers, and the work spanned many topics, from basic research and theoretical foundations to practical applications in industrial settings. It became apparent that from a mathematical viewpoint, the field has reached a certain level of maturity: Algorithmic breakthroughs were not reported, nor were they expected. Most of the talks focused on strengthening CBR with new implementation tools or extending the technology to new application areas.

The symposium opened with Alan Mackworth's (British Columbia) overview of the interplay between CBR and various logical frameworks. He emphasized that although constraint-satisfaction problems can be expressed in other logical frameworks, the relational language provides a convenient means of encoding knowledge that often invites unique opportunities for efficient processing techniques.

The discussions that followed fell into four categories: (1) extensions to commonsense reasoning, (2) parallel and distributed approaches, (3) constraint-logic-programming languages, and (4) new application areas.

The commonsense reasoning session centered on temporal reasoning issues. In particular, two approaches for combining quantitative temporal specification and qualitative specification were presented and compared. Meiri (University of California at Los Angeles [UCLA]) treats points and intervals as temporal objects of equal status, admitting both qualitative or quantitative relationships. Ladkin (International Computer Science Institute [Berkeley, California]) and Kautz (Bell Labs) maintain the qualitative and quantitative components in two separate subsystems and pro-

vide sound rules for transforming information between the two.

In the default reasoning area, Ben-Eliyahu (UCLA) and Dechter (University of California at Irvine) presented a new tractable class of default theories based on CBR mapping, and Freuder (University of New Hampshire) discussed an extension of the constraint language that expresses imprecise knowledge. In qualitative reasoning, Kuipers (University of Texas) reviewed issues in qualitative simulation and the role of constraint processing. This area concluded with Pearl's (UCLA) presentation of *causal constraint networks*, a new tractable class of constraint problem that uses the efficiency and modularity inherent to causal organizations.

Distributed and neural architectures for constraint processing received much attention. Kasif (Johns Hopkins University) surveyed the theoretical aspects of parallel computations. Attendees learned that constraint satisfaction and even arc consistency cannot be paralleled (unless the network has no cycles), meaning that they are unlikely to be solved by a polynomial number of processors in polylogarithmic time. However, linear speedup or good average parallel time is still feasible. Several distributed models were then presented, addressing whether constraint satisfaction can be achieved with neural-like computations, that is, whether constraints can be expressed as global minima of neural networks and whether these global minima can be approached by local computations.

Attendees heard answers to some of these questions: Pinkas (Washington University) showed that any set of constraints, binary or nonbinary, can be described as a Hopfield net, such that its global minima coincide with the set of solutions. Collin (Technion) and Dechter showed that even if the computation power of the neurons is upgraded to finite-state automata, the constraint problem can be solved only if one processor is made distinguished (an almost uniform model), or if the topology is a tree. Finally, Gusgen (German National Research Center) presented a uniform neural network algorithm that achieves a global solution at the expense of increasing memory requirement.

Two sessions were devoted to constraint logic programming (CLP). These languages integrate constraint satisfaction and operations research

techniques within the logic-programming paradigm. Jaffar (IBM) described the management of hard constraints in CLP systems. The idea is to delay the evaluation of nonlinear constraints until (and if) they become linear, at which point they can be solved by efficient algorithms designed for this task. Abdulla, Edpstein, Lim, and Freeman (a group from West Advanced Technology) discussed how nonlinear constraints can be managed using available packages such as MATHEMATICA. Van Hentenryck (Brown University) described improvements to arc-consistency algorithms for functional or monotone constraints and discussed their relevance to the CHIP programming language. Wilson, Borning, and Freeman-Benson (University of Washington) showed how solutions to CLPs can be obtained using hierarchical weighing of constraints and introduced control knowledge using imperative constraint programming.

Application domains were also presented. In the natural language-processing area, Haddock (Hewlett-Packard, United Kingdom) described the use of consistency algorithms to solve noun phrase reference. Kramer, Pabon, Keiroz, and Young (a group from Schlumberger) developed a polynomial algorithm for solving geometric constraint-satisfaction problems. Baykan and Fox (Carnegie Mellon University) introduced disjunctive constraints that can conveniently deal with applications of job shop scheduling and floor plan layout, and Yap (IBM) showed how the restriction site mapping in molecular biology can be expressed as dynamic constraint satisfaction.

A central issue that was repeatedly raised was the need for a standard set of large, representative, real-life benchmarks for evaluating different constraint-processing techniques.

Rina Dechter

University of California at Irvine

Implemented Knowledge Representation and Reasoning Systems

This symposium can best be thought of as two meetings: a methodological forum and a technical interchange.

On the methodological front, there was general agreement that a decrease in implemented systems was a problem for the field in the long run and that the solution was to develop a

research paradigm that encompasses the interplay between theory and implementation. Paul Cohen has made similar recommendations regarding AI in general, which he presented at the symposium and in his recent *AI Magazine* article (volume 12, number 1, spring 1991).

A central function of a paradigm in a scientific field is to determine what constitutes a *result*, that is, what is publishable. The lack of a widely agreed-on paradigm linking theory and implementation in knowledge representation and reasoning has been particularly hard on system-oriented work recently. There seems to be satisfactory agreement within the research community about what constitutes a good theoretical paper and, similarly, for the applied end of the spectrum. However, many attendees reported great confusion and frustration in the reviewing process when it came to papers describing systems that were implemented for research purposes.

Those attending had hoped to emerge with an explicit statement of a new research paradigm that addresses these problems. Unfortunately, they only identified a collection of specific pitfalls in the way that system work is currently being conducted and reported. These problems are summarized here in the form of guidelines to the authors and reviewers of papers about implemented systems.

Submerge idiosyncratic syntax: Organizing a paper around unedited transcripts of system input and output is admirable from the standpoint of "truth in advertising." Unfortunately, it often forces the reader to assimilate a large number of syntactic details that are irrelevant to the system's research contribution. Stick to commonly used syntactic conventions, except where the syntax is really the point. Verbatim transcripts can be included in an appendix if necessary.

Compare with other systems based on problem, not just technique: For example, introducing a system as "a member of the KL-ONE family" tells only part of the story. What problem is the system trying to solve? What alternative approaches have been or are being tried?

Identify the interplay with theory: It is particularly exciting to read about how an implementation was used to test some part of a theoretical proposal or how a theory was

changed based on the results of an implementation.

Provide data and examples: The role of well-chosen examples is to motivate and explain how a system works. A convincing paper must also report on results with a wider range of appropriate input.

Hypotheses must be clear and refutable: Avoid vague claims unless they can be quantified.

Identify (novel) algorithms used: To make progress in implemented systems, as researchers, we must improve our skill at sharing solutions to common subproblems. Finding the appropriate abstraction level at which to define a subproblem and a corresponding algorithm requires much effort and judgment. Compared to the more traditional parts of computer science, implementers of knowledge representation and reasoning systems have been relatively unsophisticated in this area. The most effective level for sharing implementation techniques is more abstract than the actual code in some particular system but has more details than just a complexity proof.

During the technical interchange, four themes emerged (see system descriptions in the *SIGART Bulletin* [volume 2, number 3, June 1991]):

The first theme was that almost all the implementers were concerned with the related issues of hybrid architecture and extensibility. The most ambitious goal here is to develop a general framework of interface specifications such that knowledge representation and reasoning facilities can be extended and combined with certain guarantees on overall behavior. The most interesting current approach toward this goal that was discussed is the use of object-oriented programming to define protocols for common reasoning functions.

The second theme was living with incompleteness. Most implemented systems use incomplete reasoning procedures on fairly expressive formalisms (rather than limiting expressiveness to obtain tractable complete procedures). The key challenge here is to develop principles that make the behavior of such systems understandable to their users. A promising notion toward this end is Socratic completeness: Roughly speaking, it should always be possible to construct a sequence of preparatory questions that lead the system to the correct answer to a given question.

A third theme was the need to pay

more attention to the content of knowledge representation, including high-level content and domain-specific knowledge. Speakers representing domains such as medicine, natural language, and software engineering complained that basic phenomena in these domains were still far from representable with current technology.

The fourth theme was testing and benchmarking. Although everyone agreed on the desirability of principled and sharable test data for knowledge representation and reasoning systems, almost all the participating systems were surviving with only ad hoc test suites. As an indication of the difficulties in this area, it was interesting to hear Hans-Jurgen Profitlich's (DFKI, Germany) report on the problems encountered trying to run the same test data through seemingly similar systems (seven derivatives of KL-ONE).

In summary, the symposium ended with the hope of seeing more good papers about implemented systems in the future and of having a better understanding of how to get there.

Charles Rich
Mitsubishi Electric Research
Laboratories

Integrated Intelligent Architectures

Although one of the central goals of AI is to develop artificial agents that embody all the components of intelligence, few attempts have been made to build systems that integrate multiple components, such as planning, knowledge representation, learning, vision, robotics, natural language, and interaction with other agents. The purpose of this symposium was to bring together researchers that are actively attempting to create architectures that support many of these components. Applications for attendance were received from over 100 researchers, and 45 were invited to attend, representing 35 different architectures.

The primary goal of the symposium was to help form a community that directly addresses the issues involved in creating architectures that support integration. To focus the symposium, attendees concentrated on the inter-component constraints, interactions, and synergies that arise from combining components. A secondary goal

was to get an indication of the state of the field, which was achieved by requesting that all participants address a set of questions concerning integration in their architectures, including background influences; architectural components; task characteristics that led to integration; a comparison with other systems; and an analysis of the architecture along the following dimensions: generality, versatility, rationality, programmability, ability to learn, ability to dynamically change tasks, ability to scale up language problems, reactivity, efficiency, and psychological or neuroscientific validity. Surprisingly, the participants responded to our challenge and the resulting proceedings of the symposium will be published in the *SIGART Newsletter* this year.

The represented architectures focused primarily on AI, although 11 architectures had some psychological motivations; none of them were connectionist. Overall, there is currently little coherence in the mechanisms, theory, methodology, or domains. Most of the architectures still strongly show their origins in a single component (or small number of components)—for example, as a weak-method problem solver, a temporal planner, a knowledge representation system, or a natural language-understanding system—even as they try to extend to new capabilities and new domains. There was a heavy turnout of subsumption-inspired architectures attempting to go beyond reactive behavior. There were also planners integrated with learners, knowledge representation systems extended to planning, and so on. The system that came closest to embodying a complete agent appeared to be S. Vere's HOMER, which plans and carries out actions in a simulated dynamic environment based on natural language input.

The discussions covered many topics, including the trade-offs between modular architectures and component integration; the appropriateness and utility of toy, simulated real, and real domains; and the methodology and metrics that are appropriate for evaluating systems with multiple capabilities. One characterization that came up of what the field was striving to achieve, as well as avoid, was *integration heaven* (where the components work together synergistically to generate multiplicative improvement in capabilities) and *integration hell* (where integration pro-

duces a multiplicative increase in the complexity of the architecture). In reviewing the symposium, everyone achieved a bit of heaven but not without avoiding a bit of hell.

John E. Laird

University of Michigan

Paul S. Rosenbloom

University of Southern California

Logical Formalizations of Commonsense Reasoning

The automation of commonsense reasoning is a primary goal of AI research. Commonsense reasoning, although simple for humans, is notoriously difficult to automate. AI has made much progress with seemingly difficult problems such as chess playing and medical diagnosis but remains stymied by such simple problems as blocks world planning and plausible reasoning. Part of the problem seems to be that AI researchers don't have any explicit formalization of the commonsense reasoning tasks that seem so trivial to us. Thus, designing and studying logical formalizations of commonsense reasoning domains is crucial for AI research. It forces us to make explicit the often implicit assumptions underlying commonsense reasoning, and it allows us to evaluate amorphous concepts within the rigorous framework of mathematical logic.

AI researchers, most notably John McCarthy and Pat Hayes, have noted the need for studying logical formalizations of commonsense knowledge for more than two decades. Nevertheless, this line of research has received little direct attention. Much of the work done by logic-oriented scientists is metadomain research; that is, the theories address some problem of commonsense reasoning but at an overly abstract level. Such theories can't be used to solve any particular commonsense reasoning problem in a specific realistic domain, even if they yield insights for general issues of commonsense reasoning. On the flip side, work that has been done in particular domains of commonsense reasoning typically has little theoretical content, is often ad hoc, and is not subject to rigorous analysis. In contrast to these approaches, this symposium focused on rigorous, theoretical work in specific, concrete domains.

The symposium addressed several crucial issues related to logical formal-

izations in commonsense reasoning: the presentation of logical formalizations of new domains of commonsense reasoning and of improved formalizations for previously researched domains, the discussion of general frameworks or languages suitable for commonsense reasoning, and the evaluation of existing ontologies and theories for commonsense reasoning.

Most of the papers presented formalizations of previously unexplored areas of commonsense reasoning or improved on old theories. Researchers on the MCC CYC project (which is aimed at creating an encyclopedic knowledge base of commonsense reasoning) were particularly visible. Dexter and Wanda Pratt presented a formalization of a simple theory of money, Keith Goolsbey and R. Guha presented a theory of concurrent actions, Karen Pittman presented a theory of information, and R. Guha and Doug Lenat presented an approach to reasoning using counterfactuals. Other researchers who presented theories in little studied areas of commonsense reasoning included Ernie Davis, who spoke about the kinematics of cutting solid objects, and Stephen Kaufman, who outlined a formal theory of spatiotemporal reasoning. Kate Sanders presented a theory of emotions; Jennifer Elgot-Drapkin suggested a solution to the logical omniscience problem; and Yoav Shoham, Anton Schwartz, and Becky Thomas proposed a model for the mental state of an agent. Some papers addressed more general topics that are central to commonsense reasoning, such as plausible reasoning (Jim Delgrande and Ernie Davis), reasoning with incomplete information (Michael Gelfond), and causation (Charles Elkan). A minority of papers addressed traditional difficult problems in commonsense reasoning. Ray Reiter, Murray Shanahan, and Fangzhen Lin and Yoav Shoham addressed various aspects of the frame problem. Finally, a number of papers presented general frameworks for various aspects of commonsense reasoning. These frameworks included a proposal for extending the situation calculus so that it can accommodate concurrent actions and partially specified actions (Michael Gelfond, Vladimir Lifschitz, and Arkady Rabinov) and a formal system allowing the use of multiple distinct languages for different reasoning purposes (Giunchiglia).

A highlight of the symposium was the panel on theory evaluation in AI. Although many AI researchers generate theories, there has been little effort to rigorously evaluate existing theories, compare new theories with old, and judge how well a theory solves a set of problems. Several of the panelists argued the need to establish a set of criteria by which AI theories should be judged, and there were suggestions and discussion of several concrete evaluation criteria. What constitutes a reasonable set of benchmark problems was also discussed at length. This panel was of particular interest to researchers interested in commonsense reasoning, but in fact, the issues addressed are important ones for the AI community in general: How can we tell if this theory (or system) is really good? What makes this theory better than existing theories? All AI researchers should be prepared to answer these questions when they present new work.

During the open discussion on future directions for commonsense reasoning, there was evident excitement for the sort of specific domain theory building that the people in the CYC project are doing. This excitement was tempered by a recognition of the need to integrate this work with existing AI theory and aim toward building problem-driven theories. There was a general consensus that meetings of this sort were too few. For many, this gathering represented the only chance to attend a symposium where the main focus was on the presentation of formal, concrete theories in specific aspects of commonsense reasoning. Papers of this kind have often been presented at other conferences, but the focus on commonsense reasoning frequently gets lost. Participants plan to repeat this positive experience in about a year and a half and hope that the enthusiasm for the field of commonsense reasoning continues to grow.

Leora Morgenstern

IBM T. J. Watson Research

Machine Learning of Natural Language and Ontology

People begin their lives without the ability to speak any natural language and are able, in a few short years, to develop a linguistic competence that enables them to function as a writer, scholar, politician, whatever they

choose to become. They might, in fact, learn to communicate in several natural languages. These remarkable phenomena of language learning have amazed most of us at one time or another, and it is only natural that we have tried to use computers to study or even duplicate them—with only partial success to report at this date.

The AAAI Spring Symposium on Machine Learning of Natural Language and Ontology (MLNLO) provided an opportunity to get together and discuss the partial successes and the research challenges that lie ahead. It was a rare opportunity because the work has tended to be reported in fragments, a thesis here or there, a paper at an AI or computational linguistics conference, another at a psychology or linguistics or child language conference or in a philosophy journal. The field is naturally highly multidisciplinary, and the interested researchers all speak their own languages—not just natural languages but specialized disciplinary dialects laden with the theoretical constructs and assumptions of each discipline. Thus, this symposium provided a forum for useful interchange of ideas.

“Learning of natural language” is a simple-sounding phrase that covers a number of phenomena. On the one hand, there are various aspects of language to be learned, such as the sounds that are significant to a particular language (phonology), words (lexicon) and their variations (morphology), the structure of meaningful utterances (syntax), and meaning and its relation to the lexicon and syntactic structure (semantics). On the other hand, there are the different components of learning: inducing the data to be learned from raw linguistic and nonlinguistic data; somehow codifying these data into an internalized, structured system that can be used in an automatic manner; and generalizing to be able to deal with new input never heard before and produce new output never uttered before. The learning of *ontology*, the understanding of what exists in the world, is closely linked with the learning of language.

At the symposium, 50 participants discussed contributions in all these areas, with 20 full-length presentations and a similar number of “advertising spots” that allowed virtually all groups some air time. It should also be mentioned that a parallel symposium focused on connectionist natural language processing (CNLP) and

that nontraditional computing has clearly exerted its influence on the field of language learning. Not only were a number of applications of connectionist and genetic techniques presented, but a joint final session was held with CLNP.

However, the main efforts are still closely linked to contemporary AI and linguistic theory. The field is beginning to attack various practical applications in areas where the knowledge is rich enough to allow modest learning, and it is providing increasing challenge and support to psycholinguistics and linguistics research. In this respect, participants spent time focused on a number of special topics, such as the extent to which language mechanisms are language specific and linguistic properties are innate, the conditions under which it is formally possible to learn a language, the recognition of ungrammatical sentences, the development of the ability to use metaphors, the modeling of second-language learning, and the question of how lexical symbols become grounded in reality. The treatments presented complemented cognitive theory with computational implementation.

At the end of the symposium, we took time before the joint CNLP panel to review the value of the symposium and look to the future. It was resolved that we instigate a regular program of MLNLO events; a newsletter; resource sharing (software, texts, and so on); and further symposia, workshops, and conferences. The first such event was a one-day workshop on natural language learning to be held at the 1991 International Joint Conferences on Artificial Intelligence in Sydney, Australia, on August 25. As befitted its shorter length, this workshop had a tighter focus, with a major goal being an analysis of proposed language-learning models to allow comparing and contrasting of the theoretical perspectives and the hypotheses embodied; the implementation techniques and learning algorithms; and the implications of the virtues, failings, and results of particular implementations and modeling experiments.

The symposium participants also felt that the working notes of the MLNLO symposium were a landmark volume worthy of further distribution. Thus, the working notes will immediately be made available (through the German AI Institute in

Kaiserslautern [DFKI D-91-09]) to a wider audience in the form of a technical report and an edited book. These publications will allow the expanded presentation of selected papers and, perhaps, additional invited papers from some who could not attend. Information can be obtained from powers@informatik.uni-kl.de or from reeker@ida.org.

David Powers

University of Kaiserslautern

Larry Reeker

Institute of Defense Analyses

Principles of Diagnosis

(continued from page 23)

Constraints: A Language for Expressing Almost-Hierarchical Descriptions. *Artificial Intelligence* 14(1): 1–40.

Szolovits, P., and Pauker, S. G. 1978. Categorical and Probabilistic Reasoning in Medical Diagnosis. *Artificial Intelligence* 11:115–144.

Weld, D. S., and de Kleer, J. 1990. *Readings in Qualitative Reasoning about Physical Systems*. San Mateo, Calif.: Morgan Kaufmann.

Williams, B. C. 1986. Doing Time: Putting Qualitative Reasoning on Firmer Ground. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, 105–112. Menlo Park, Calif.: American Association for Artificial Intelligence. Also in 1990. *Readings in Qualitative Reasoning about Physical Systems*, eds. D. Weld and J. de Kleer, 353–360. San Mateo, Calif.: Morgan Kaufmann.

Wu, T. D. 1991. Domain Structure and the Complexity of Diagnostic Problem Solving. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, 855–861. Menlo Park, Calif.: American Association for Artificial Intelligence.

Wu, T. D. 1990. Efficient Diagnosis of Multiple Disorders Based on a Symptom Clustering Approach. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 357–364. Menlo Park, Calif.: American Association for Artificial Intelligence.

Zadrozny, W. 1990. The Logic of Abduction. In *Working Notes of the First International Workshop on Principles of Diagnosis*, ed. W. C. Hamscher, 8–17, Technical Report 14, Price Waterhouse Technology Centre, Menlo Park, California.



Walter Hamscher, who organized and chaired Dx-90, earned his Ph.D. in artificial intelligence from the Department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology in August 1988. He is currently a research manager in the Systems Research Group at Price Waterhouse Technology Centre, where he and his colleagues are developing a model-based approach to a variety of tasks in auditing and financial analysis.