

Introduction to the Special Issue on Question Answering

David Gunning, Vinay K. Chaudhri, and Chris Welty

■ *This special issue of AI Magazine presents six articles on some of the most interesting question-answering systems in development today. Included are articles on Vulcan's Project Halo, Cyc's Semantic Research Assistant, IBM's Watson, True Knowledge, and the University of Washington's TextRunner.*

Starting from the Turing test, question answering (QA) has often been thought of and used as a direct way of observing and measuring intelligent behavior in machines. Even though AI has diversified much beyond the notion of intelligent behavior proposed in the Turing test, QA remains a fundamental capability needed by a large class of systems. The QA problem extends beyond AI systems to many analytical tasks that involve gathering, correlating, and analyzing information in ways that can naturally be formulated as questions. Ultimately, questions are an interface to systems that provide such analytic capabilities, and the need to provide this interface has increased dramatically over the past decade with the explosion of information available in digital form.

The suitability and simplicity of the question-asking interface is so obvious that, since the earliest days of science fiction, nearly every story in which computers are present assumes them to be talking machines that effortlessly understand and answer any conceivable question. Yet, the best we experience in reality today is the clever use of key-word search to find one or more related documents. We expect intelligent machines to be able to answer questions — to provide specific answers to specific questions by understanding, synthesizing, and reasoning about the underlying data, knowledge, and text documents. We also expect it to be able to provide user-specific explanations or justifications relating the context of the question to what the user currently knows, as well as informing the user of its confidence, especially in cases where the confidence in an answer is low.

For 50 years, AI researchers have chipped away at answering questions with logic, programs, expert rules, statistics, information retrieval, and so on. More recently, with advances in knowledge-based systems, natural language understanding, machine learning, and text understanding across the web, we may be on the threshold of finding combinations of these techniques that achieve a convincing and useful question-answering capability that is radically different from what is available in the market today.

	Cyc SRA	Project Halo	IBM Watson	True Knowledge	TextRunner
Domain	Closed	Closed	Mixed	Open	Open
Knowledge Representation	Higher Order Logic	Description logic plus Horn Rules	Mixed	Triples	Triples of text
Knowledge Acquisition	Manual	Manual	Mixed	Mixed	Automatic
Question Form	Interactive Natural Language	Controlled Natural Language	Full Natural Language	Question Templates	Full Natural Language
Reasoning	Deductive	Deductive	Mixed	Production Systems	Matching
Explanation	Yes	Yes	Provides Confidence	Yes	No

Table 1. System Differences.

In this special issue, we have assembled articles on some of the most interesting question-answering systems in development today. Included are articles on Project Halo (Vulcan's effort to create a scientific knowledge base that can answer questions and explain the answers), the Semantic Research Assistant (one of Cyc's most interesting applications that answers ad hoc queries for clinical researchers at the Cleveland Clinic), Watson (IBM's ambitious effort to play —and win — the game of Jeopardy against human contestants), True Knowledge (one of the first and best-known question answering systems on the web today), and TextRunner (University of Washington's clever use of statistical natural language techniques to answer questions across the open web).

There are of course many others. We have tried to select a set that represents the state of the art in implemented question-answering systems today and that covers a range of approaches being pursued. On one hand, we have systems like Cyc SRA that use a rich knowledge representation but cover a closed, manually constructed knowledge domain; on the other, we have systems like TextRunner that use a much simpler representation but cover the completely open domain of the web. We have AURA, Watson, and True Knowledge in between, with Watson representing perhaps the most hybrid mixture of structured and unstructured techniques.

The systems differ on many dimensions. The domain may be open (the

entire web) or closed (a section of a science textbook). The knowledge representation may be richly structured (first-order or higher-order logic) or simply structured (triples of text fragments). Knowledge acquisition may be fully automatic (by information extraction across the web) or fully manual (by entry from domain experts). The form of the questions that can be asked may be simple (in a controlled language) or complex (open to full natural language). The reasoning performed to answer the question may be simple (matching the question to an underlying knowledge structure) or complex (involving multiple step inference). Finally, the system may or may not provide an explanation of its answer.

Table 1 gives a summary of how these systems compare on these dimensions. This characterization is oversimplified at best. Each system has a rich story to tell regarding each of these dimensions. We have tried to briefly sketch some of the predominant characteristics of each system to give the reader a starting context.

The systems we've chosen here demonstrate that solving the question-answering problem requires breakthroughs in knowledge representation, knowledge acquisition, deductive inference, natural language processing, and machine learning — the pillars of artificial intelligence — and will require the integration of these techniques in new and interesting ways. It has produced, and we believe will continue to produce, significant

advances in AI that will generalize to other problems and the overall quest of building intelligent machines.

AI grand challenges often seem unreachable — before they are solved. Playing chess, understanding speech, and even driving autonomously all seemed impossible at one time, but now appear much less mysterious. Achieving a system that genuinely answers questions may now be on the threshold of becoming a reality. It may very well be one of these systems, or certainly techniques pioneered by these systems, that will make that dream possible.

David Gunning is a senior research program manager at Vulcan Inc., leading the AURA and HaloBook developments. Prior to joining Vulcan, he served as a program manager at the US Defense Advanced Research Projects Agency (DARPA), where he led numerous AI projects including the Personalized Assistant that Learns (PAL) and the Command Post of the Future (CPOF).

Vinay K. Chaudhri (Vinay.Chaudhri@sri.com) received his Ph.D. from the University of Toronto. He is currently a program director for ontology management in the Artificial Intelligence Center at SRI International. His research focuses on knowledge acquisition, ontologies, and deductive question answering.

Chris Welty is a research scientist at the IBM Thomas J. Watson Research Center. His background is primarily in knowledge representation and reasoning. Welty's current research focus is on hybridization of machine learning, natural language processing, and knowledge representation and reasoning in building AI systems.