Praveen K. Paritosh

Statement of Research Interests

I am interested in understanding the nature of human commonsense reasoning by building computational models of it. Commonsense reasoning provides a solution to overcoming brittleness in Artificial Intelligence (AI) systems. It is also a necessary prerequisite to creating useful software that collaborates with human partners to accomplish tasks like damage control assessment, operations planning, sifting through on-line information for relevant data, teaching and tutoring, and developing complex scientific and engineering models. It also provides a very interesting vantage point to explore cognitive aspects of problem solving, categorization and memory.

In my research, under the guidance of Ken Forbus and Dedre Gentner, I have designed and implemented problem solving systems that demonstrate the broad coverage of a small set of commonsense reasoning strategies. I have built a program that makes back-of-the-envelope estimates when there isn't enough information for an accurate answer, for questions from both commonsense and scientific domains like "How much money is spent on healthcare in the US?" Making good estimates involves having good intuitions about quantities, or, quantity sense. Quantity sense consists of knowing what is reasonable, high and low, and knowing the important points, e.g., phase transitions. I have developed a cognitive theory of how people acquire quantity sense, which has been implemented as a computational model that makes novel but psychologically-testable predictions. The next two sections describe my work, and then I talk about future research directions.

1. Back of the Envelope Reasoning: A Solution to the Brittleness Problem

Brittleness is a key problem for most AI programs, and perhaps software in general. The two common manifestations of brittleness are: 1) the software cannot find an answer, because of gaps in its knowledge base or because of lack of computational resources required, and 2) the software comes up with an unreasonable answer, possibly because of inaccuracies in its knowledge base (garbage in/garbage out). For instance, in an evaluation of question-answering programs that mine text for answers[1], one program came up with 360 tons as the amount of folic acid that an expectant mother should have per day, and 14 feet as the diameter of the earth! A solution to the brittleness problem is to equip programs with commonsense reasoning, specifically, the ability to make educated guesses as a fallback mechanism when resources (knowledge and/or computation) are lacking. This can also be used for sanity checking of other reasoning processes and for fact checking knowledge acquired from noisy sources like the web.

Back of the envelope (BotE) reasoning involves generating quantitative answers in situations where exact data and models are unavailable, and where available data is often incomplete and/or inconsistent. Such reasoning is a key component of commonsense reasoning about everyday physical situations. I have built BotE-Solver, a general-purpose problem solving framework that uses *estimation strategies* and the ResearchCyc knowledge base. The power of BotE-Solver comes from its strategies that enable it to come up with an answer even when none can be found using standard methods. A strategy transforms a given question into other, possibly easier questions. A key contribution of this work is that a core set of seven strategies provides broad coverage, and is arguably the complete set of back of the envelope problem solving strategies. There is twofold support for this hypothesis: 1) an empirical analysis of all problems ($n=44$) on Force and Pressure, Rotation and Mechanics, Heat, and Astronomy from Clifford Swartz's book, "Back-of-the-Envelope Physics," and 2) an analysis of problems solved by BotE-Solver.

---

[1] The question is from TREC9, and this was reported in the IBM TJ Watson AQUAINT Briefing.

A BotE question asks for an estimate of a quantity for some object, which can be abstractly stated as (Q O ?V), where Q is the quantity, O the object, and ?V the unknown value. This suggests three syntactic transformations, namely, transforming the object, quantity, or both. An example of an object-based strategy is mereology, i.e., decomposing an object into its sub-parts. If the quantity in question is an extensive quantity, like weight, its value for the whole is sum of parts; while for an intensive quantity, like density, its value for the whole is the weighted average of the values for the parts. An example of a quantity-based strategy is using a physical law like $F=m*a$. Applying conservation laws or balances are examples of system-based strategies. When asked a question, BotE-Solver first tries to see if the answer is available in the knowledge base. Failing that, it tries to find similar examples for which answers are available. This is the analogy strategy, an important object-based strategy. If no analogues are found, then other applicable strategies are recursively applied.

2. A Cognitive Model of Quantity Sense

In order to use analogies to make numeric estimates, the analogical matching algorithms should be sensitive to quantities in the first place. Most models of similarity do not adequately handle numeric properties – either ad hoc similarity metrics such as Euclidean distance are used, or the quantities are completely ignored in the matching and retrieval processes. My theory of quantity sense claims that learning about quantities consists of building qualitative representations for them. This learning process can be summarized as *symbolization by comparison*. Comparison helps us notice and extract the scale of values of quantities and we create symbolizations that name points and intervals on this scale. These symbolizations must make two kinds of distinctions: *distributional*, those that denote changes of quantity, e.g., large and small; and *structural*, those that denote changes of quality, e.g., boiling point and poverty line. This has been implemented as a computational model, CARVE. It takes as input a set of examples represented in predicate calculus, and finds and generates qualitative representations corresponding to the distributional and structural partitions. This has strong implications for models of similarity, retrieval and generalization. My theory takes a radically different approach to the problem of incorporating quantities in similarity models by proposing that the solution lies in better representations, not in the similarity metrics. I present existing evidence from psychology and linguistics to support the theory. A functional validation of the theory is that the representations generated by CARVE help BotE-Solver generate more accurate estimates.

3. Future Work

3.1 Integrating Problem Solving and Question Answering

Some immediate applications of my work on back-of-the-envelope reasoning are to build tools for everyday numeracy support, for example, *number-checker* (and *number-explainer*) like the spell-checker that can alert and provide explanations when a number does not make sense, or a search engine that is geared towards finding and generating numerical estimates. Imagine being able to click on a number in a news article or a financial report, and being offered a back-of-the-envelope estimate showing how it makes sense, or being shown other comparable quantities that contextualize it. These techniques could be used to build both end-user tools and middleware for information extraction and knowledge acquisition, where sanity-checking is an important issue. The next research steps will be looking at meta-reasoning in constraining commonsense reasoning, and building support for natural language sources like books, newspapers and the web as knowledge bases and for interaction. Connecting to language is important in extending the scope of knowledge-based systems. Knowledge-based systems can make sophisticated inferences, however, currently there exists no knowledge base that is broad enough to perform well in the TREC question answering competitions. Integrating knowledge- and text-based system will result in answers to: 1) What kinds of reasoning can be accomplished with textual representations? And 2) How can structured knowledge be acquired from text? The current work on textual entailment is beginning to look at similar issues, and I believe by starting with a well-understood reasoning domain like back of the envelope reasoning is likely to achieve more success.

3.2 Cognitive and Educational Aspects of Quantity Sense

I would like to explore the psychological ramifications of the symbolization by comparison theory. Recent work has shown that humans along with many other animals share a cognitive infrastructure for representation of approximate numerosity[2]. However, there exists an explanatory gap between how our qualitative representations of quantity are related to this cognitive infrastructure. This line of research has educational significance as well. Quantitative literacy is a very important issue for math education, and there are many demonstrations of the lack of success of current educational methods to impart this skill at various levels from middle-school to college. For example, more than 90% of mechanical engineering seniors (100 at MIT, and 250 from five other universities) came up with estimates that were off by more than one order of magnitude for the value of energy stored in a 9-volt "transistor" battery, and responses varied by nine orders of magnitude[3]. I am interested in using the cognitive insights from my work and others to design an undergraduate-level class centered on teaching estimation skills. My research on back of the envelope reasoning provides a framework to structure the class, and this class will also serve as a laboratory to generate and explore hypotheses about human commonsense reasoning to guide my research.

3.3 Heuristic Reasoning and Organization of Knowledge

My goal is to understand how people organize their knowledge to make educated guesses, as I believe that is a major source of the flexibility and robustness of human intelligence. The body of psychological literature stemming from Tversky and Kahneman's *Heuristics and Biases* program[4] provides some suggestion about organization of knowledge and inference patterns that underlie human commonsense reasoning. On the face of it, my proposal might seem paradoxical: I want to look at the biases, flaws, errors and limitations of human reasoning to build better reasoning engines. But on the same token, human reasoning is more flexible and powerful in most situations than anything else we know about. The psychological work on human reasoning is a large collection of phenomena, usually highlighting deviations from normative/rational reasoning. There is hardly any work in AI looking at how these results could be useful for designing AI systems. A source of difficulty is that the psychological results come from various different experimental paradigms and different representational assumptions. Another difficulty lies in the knowledge- and memory-impoverished nature of many psychological studies. I am interested in organizing psychological work in a rigorous computational framework and exploring the organization of knowledge and long-term memory that makes many of the heuristics work. The motivation is not to just simulate biases, but to discover methods that make the reasoning more flexible/better most of the time. A strong constraint will be building integrated architectures for the mind: it might be trivial to build a computational model of a particular bias. I believe that this approach will lead to a deeper understanding of cognition.

---

[2] Dehaene, S. 1999. The Number Sense: How the Mind Creates Mathematics, Oxford University Press.

[3] Linder, B. 1999. Understanding Estimation and its Relation to Engineering Education, Ph.D. dissertation, Department of Mechanical Engineering, MIT

[4] Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and biases. *Science*, **185**, 1124-1130. Also, more recently, Gigerenzer, G., Todd, P. M. & the ABC Research Group. (1999). Simple Heuristics That Make Us Smart. New York: Oxford University Press.