

Turing's Dream and the Knowledge Challenge

Lenhart Schubert
University of Rochester

Abstract. There is a set of clear-cut challenges, all centering around knowledge, that have received insufficient attention in AI, and whose solution could bring the realization of Turing's dream – the dream of a machine we can talk with just like a person, and which is therefore (at least) our intellectual equal. These challenges have to do with the representation of linguistically expressible knowledge, the role of knowledge in language understanding, the use of knowledge for several sorts of commonsense reasoning, and knowledge accumulation. Concerning the last topic, I briefly present preliminary results of some of our recent efforts to extract “shallow” general knowledge about the world from large text corpora.

Hi-fidelity Representation of Linguistically Expressible Knowledge

“Language is fundamental to our ability to think; it is ‘more or less synonymous with symbolic thought.’...”

– Donald Johanson, *Sci. Am.* 1998, expounding on Ian Tattersall's *Becoming Human*

Humans acquire large amounts of knowledge about the world from verbal expressions of such knowledge, and in turn, are able to communicate much of their internalized knowledge in language. It stands to reason, therefore, that our ‘mentalese’ must match the expressive resources of language – and that artificial agents competent in language will also require internal representations no less expressive than language. My collaborators and I have argued for many years (e.g., (Schubert & Hwang 1989; 2000; Schubert 2000)) that a representation capable of supporting natural language and commonsense reasoning in machines must allow not only for predication, logical connectives, and \forall/\exists -quantifiers – the devices of classical logic – but also the following:

- **generalized quantifiers** such as *most* or *often*; e.g., “*Most of the trick-or-treaters who came to our door received several candy bars*”; though set-theoretic paraphrases are possible, inference is more straightforward if such quantifiers are directly allowed for;
- **event/situation reference**; e.g., (following the previous sentence) “*This nearly emptied our bowl of treats by 7:30*

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

pm”; note the anaphoric reference to the collection of ‘receiving events’, and the temporal modification of the ‘nearly-emptying event’;

- **modification (of predicates and sentences)**; examples are “*nearly*” in the previous sentence, verbs like “*become*” and “*pose as*” (with a predicate complement), and sentence adverbials such as “*perhaps*” and “*according to reliable sources*”; many such examples require an *intensional semantics*, i.e., one that does not reduce meaning to reference (extension) in the world;
- **reification (of predicates and sentences)**; examples are “*Domestic dogs evolved from wolves and African wild dogs*” (nominal predicate reification), “*Composing music requires problem solving, but is deeply gratifying*” (VP predicate reification), “*That Turing was brilliant is beyond doubt*” (proposition reification), “*For Turing to make a mistake was unusual*” (event-type reification), and “*Who murdered Montague remains a mystery*” (question reification); again intensionality is involved in many such examples;
- **metric/comparative attributes**; e.g., “*Thousands of men were employed for 20 years to build the Great Pyramid, which rose to a height of 485 feet*”; “*The frame of the bed is wider than the door is high*”;
- **uncertainty and genericity**; e.g., “*If John receives a job offer from North Central Positronics, he will probably accept it*”; “*Elementary school children are usually bused to school*”; “*Dogs bark*”;
- **metalinguistic capabilities**; e.g., “*What is the sum of binary numbers 111 and 1, expressed as a binary number?*”; “*Can you finish the opening sentence of Hamlet's soliloquy, ‘To be or not to be, ...?’*”;

It is possible to contrive FOPC paraphrases by hand for some of the above examples, but doing so automatically is deeply problematic, and probably impossible in some cases (e.g., for some examples of modification, reification, genericity, and metalanguage). Our Episodic Logic (EL) representation (Schubert & Hwang 2000) covers many of the above constructs and implements them in the EPILOG system (Schaeffer *et al.* 1993; Schubert & Hwang 2000). Also (Fox & Lapin 2004; 2005) provide a two-level representation (PTCT –

Property Theory with Curry Typing) wherein propositions and other intensional entities are represented as terms, and these are embedded in a first-order language with a predicate **True**. They provide a tableau proof theory for the language, and their focus is on the logic, more than on its exact connection to various linguistic phenomena, or on computational issues.

Much additional work is needed to develop semantics and practical inference methods, including various types of nondeductive inference, for KR languages with the above sort of expressive range. Worries about tractability of such languages are quite misguided and may have hindered progress on NLU. Promoting research into tractable inference is one thing, but inveighing against formalisms that allow intractable inference is like inveighing against programming languages that allow recursion or lambda-abstraction.

The Role of Knowledge in Language Understanding

One of the earliest and most important lessons learned in AI research on NLU is that language understanding is knowledge-intensive. It requires simultaneous deployment of a vast amount of pattern-like knowledge (about preferred syntactic, semantic, and discourse patterns) and propositional knowledge (about lexical meanings, discourse, and most of all, the world).

So far, unfortunately, genuine knowledge-dependent understanding has been achieved only in micro-domains, where the vocabulary, syntax, and content of the discourse are so limited that adequate interpretation can be achieved with domain-dependent rules. For example, in domains like making hotel or flight reservations or performing other commercial transactions, or assembling, trouble-shooting, or repairing simple devices, it may be quite feasible to use specialized rules that look for particular types of word/phrasal patterns and use these to hypothesize slot-fillers for templates that are geared towards the micro-domain.

But such an approach requires a re-engineering effort for each new micro-domain, and leads to brittleness at domain boundaries and difficulties in dealing with novel combinations of situations. Moreover, making connections between the linguistic input and any known templates or scripts may require multiple knowledge-based inferences that cannot be anticipated by heuristic mapping rules. Consider, for instance, the following brief story:

John was driving on a busy road when he came up behind a flatbed truck carrying concrete sewer-pipe sections; oddly, a man who appeared to be a hitchhiker with some object in his hand was crouching on the pipes. As John pulled closer, he suddenly realized that the man was using metal shears to cut the straps securing the pipes...

Human readers effortlessly foresee the potential consequences of this scenario, enabling them to comprehend a continuation such as “*John braked hard and swerved onto the shoulder*”. For general understanding, logical forms derived from surface structure must mesh with general and contextual knowledge in a way that allows inference,

where the mode of inference may be largely conjectural and strongly dependent on syntactic, contextual, and semantic biases.

A promising abduction-like framework has been suggested by Hobbs *et al.* (Hobbs *et al.* 1993), and this has been successfully applied in selecting best answers in text-based question answering (Moldovan *et al.* 2000). The idea is to seek an ‘explanation’ for the truth of each sentence of a text, in the form of a minimal-cost proof of the sentence from prior knowledge. An important advantage of the approach is that it allows assumption-making, at a ‘cost’, in deriving an explanation for a sentence; thus it allows for the incompleteness of background and contextual knowledge. However, I believe that the FOPC-based, purely conjunctive logical forms used so far in that approach are not general enough; that beliefs, goals, obligations, etc., of discourse participants as well as individuals under discussion need to be taken into account; that inference needs to be much more data-driven, generating speculative predictions and explanations from new inputs; and that speculative inferences generated from one sentence create a context for the next sentence, powerfully affecting its interpretation. For example, the ‘explanation’ for the sentence “*John began to run*” is quite different for the two possible preceding sentences “*John heard steps behind him*”, and “*John looked at his watch*”.

Commonsense Reasoning and Planning

The research directions indicated by the inference requirements of a Turing-like dialogue agent, and human-level reasoning in general, include

- **Reasoning and planning in an enriched logical syntax** (as above), and with selective knowledge-access strategies that allow scaling up to very large knowledge bases; our EPILOG inference engine is an attempt in that direction (Schubert & Hwang 2000; Schaeffer *et al.* 1993), allowing goal-driven and input-driven inference (but not planning) in EL, and accessing formulas via ‘key predicates’ together with arguments or argument types; there is a somewhat similar retrieval mechanism in MAC/FAC, based on predicate and function occurrence counts (Forbus, Gentner, & Law 1994).
- **Several varieties of uncertain inference:** *Causal/explanatory inference, abduction:* These (related) inference modes involve variable degrees of uncertainty, and as such run into the problem of ‘where the numbers come from’ in uncertain inference; in general our ‘statistical experience’ radically underdetermines the distributions that seem to be required for well-founded uncertain inference. Perhaps general logics incorporating Bayesian-network-like inference (and thus, strong causal independence assumptions) will emerge as suitable tools (e.g., (Poole 1993; Ngo & Haddawy 1996; Pfeffer 2000; Halpern 2003; Schubert 2004; Richardson & Domingos 2006) provide some push in that direction).

Analogy: analogies powerfully and pervasively shape our expectations about the world. Our analogical tendencies

are perhaps best appreciated from an outlandish example that lies outside common experience and knowledge:

You are exploring a distant planet and encounter a beach-ball-sized, silvery, globular entity floating two feet off the ground; it splits open, releasing a smaller blue globe that drifts two feet upward, and bursts, spattering blue droplets over you; these produce no apparent ill effect, but you are disconcerted, and hastily depart the scene – but minutes later you encounter a beach-ball-sized, silvery, globular entity floating two feet off the ground; it splits open, releasing a smaller red globe ...

The reader presumably has certain expectations about what might happen next. The point is that a *single* conjunction of objects or events often suffices for analogical transfer to similar conjunctions of objects or events, allowing prediction of as-yet unconfirmed aspects of the conjunction. A noteworthy effort to incorporate this type of analogical inference into an NLU system is the Learning Reader project at Northwestern University (Ureel *et al.* 2005), making use of techniques developed as part of the MAC/FAC system (Forbus, Gentner, & Law 1994) and SEQL (Kuehne *et al.* 2000). Many questions remain, however – what makes a conjunction of objects or events (and certain of their properties) a plausible candidate for analogy, distinguishing it from accidental, nonpredictive conjunctions? How should multiple analogous conjunctions be combined into more abstract, more general ‘prototypes’? (Here a line of work on concept formation going back at least to (Winston 1975) is relevant.) How can we generalize to non-ground, logically complex descriptions of sets of objects or events? How does analogy-based and prototype-based inference relate to inference based on generic propositions of the sort mentioned earlier?

‘Distillation’ of the stream of experience: consider simple claims such as “*I did a lot of shopping and cleaning today*”, or “*The novella is about an aging, destitute Cuban fisherman who catches a huge marlin in the Gulf Stream but loses it to sharks*”; obviously, the first claim is an abstraction from a welter of remembered experiences and activities in the course of a day, and the latter is similarly an abstraction from a large number of remembered details in the novella. This important type of abstraction process has been largely ignored in AI and NLU, notwithstanding the burgeoning field of text summarization.

Concept and theory formation: Here I am not referring to learning to classify instances of predetermined categories, but rather to predicate and theory ‘invention’, based on recurrent constellations of objects, structures, or events – where the usefulness of the abstracted predicates and theories lies in allowing more concise description and effective prediction of some domain. Some work in this direction has been done in automated mathematical theorizing, invariant discovery in planning, and relational data mining.

- **Continuous, utility-guided planning and action;** to be human-like, a conversational agent needs to exhibit initia-

tive and pursue adoptive as well as self-generated goals systematically, continually updating its current overall plan, with allowance for uncertainty and trade-offs; work on continuous planning, for instance at JPL in Pasadena, can be considered a start in that direction (Chien *et al.* 1999).

- **Specialized inference** for taxonomies, times, parts, geometry/ motion, etc.; e.g., the following are trivial for people but not necessarily for general reasoners: “*Is a sequoia a living thing?*”; “*Is a squirrel a kind of monkey?*”; “*What comes first in ‘Hamlet’, the death of Polonius or the death of Laertes?*”; “*Can you put an open umbrella into the trunk of a standard-size sedan?*”. Our EPILOG system incorporates a dozen specialists through a central interface – though not the sort of imagistic reasoner the last example hints at.
- **Ordinary reasoning involving mathematical induction;** D. McAllester has posed a number of intuitively trivial problems that are hard for machines (McAllester 1991); An example is whether a king on an otherwise empty chessboard can reach all squares; McAllester offers some inductive inference rules, but fully adequate methods remain elusive.
- **Meta-inference, introspective inference and self-awareness;** this emerging area (e.g., (McCarthy 1999; Cox 2005; Schubert 2005)) is essential to the construction of agents that display ‘presence of mind’ in a dialogue, with awareness of what has transpired so far, what is expected of it, what it can and cannot do, and what it knows and does not know.

Knowledge accumulation

At 1.2 million propositions, the Cyc project has achieved the largest compilation so far of general world knowledge (Lenat 1995). But arguably the requirements for human-level understanding are one or two orders of magnitude higher. Furthermore, NLU seems to require many millions of items of preferential pattern-like knowledge not covered by Cyc, and even much of the world knowledge needs to be recast in generic form (e.g., domestic dogs are *generally* friendly, airplanes *often* depart and land behind schedule, etc.)

An extremely attractive potential source for knowledge accumulation is the plethora of electronically encoded lexicons, encyclopedias, and other books that are available nowadays, and billions of miscellaneous web documents. However, there is a bootstrap problem: reading and understanding these materials largely presupposes the general knowledge we are trying to acquire.

As a way of making a start, we have been working at the University of Rochester on acquiring shallow, pattern-like knowledge from large text corpora, based on the idea that the patterns of predication and modification in language tend to reflect frequently observed relationships in the world (Schubert 2002; Schubert & Tong 2003). This approach has yielded over 100,000 general propositions from the Brown corpus and over 2 million from the British National corpus. Examples like “A PERSON MAY BELIEVE A PROPOSI-

TION”, “BILLS MAY BE APPROVED BY COMMITTEES”, “A US-STATE MAY HAVE HIGH SCHOOLS”, and “CHILDREN MAY LIVE WITH RELATIVES” provide some idea of the sorts of shallow knowledge obtained.

The extraction method is based on compositional derivation of logical forms (in EL) from parsed sentences, abstraction of arguments (e.g., abstracting *IBM* to a *company*, or *she* to a *female-individual*), and dropping of lower-level modifiers at higher structural levels. In general, several propositions may be derived from a sentence, based on sub-clauses and modifiers. For example, given the Treebank bracketing of the sentence

*Rilly or Glendora had entered her room while she slept,
bringing back her washed clothes,*

our programs produces a set of 5 logical forms, which are then verbalized as “A FEMALE-INDIVIDUAL MAY ENTER A ROOM”, “A FEMALE-INDIVIDUAL MAY HAVE A ROOM”, “A FEMALE-INDIVIDUAL MAY SLEEP”, “A FEMALE-INDIVIDUAL MAY HAVE CLOTHES”, and “CLOTHES CAN BE WASHED”. More recently work along similar lines has been done, for example, at Boeing (Clark, Harrison, & Thompson 2003) and Univ. of Washington (in the context of the KnowItAll project, <http://www.cs.washington.edu/research/knowitall/>). The shallow, pattern-like knowledge extracted in such work seems promising for guiding parsers and logical form generators, but at present it falls short of providing the kind of conditional knowledge that can be used for reasoning.

Conclusion

Obviously, the challenges that still confront the achievement of Turing’s dream are numerous and difficult. But over the last half-century, these challenges have turned from mysteries into problems. We can enumerate the problems, and point to significant progress on all of them. They center around knowledge and its relation to language, its use for inference and its accumulation. The time is ripe for moving from oversimplified representations to ones that match the richness of language, for moving from mere deduction to more speculative modes of inference (within an expressive representation), and for more fully exploiting online text resources for knowledge acquisition.

Acknowledgement

This work was supported by NSF grant IIS-0328849.

References

- Chien, S.; Knight, R.; Stechert, A.; Sherwood, R.; and Rabideau, G. 1999. Integrated planning and execution for autonomous spacecraft. In *Proc. of the IEEE Aerospace Conference (IAC)*.
- Clark, P.; Harrison, P.; and Thompson, J. 2003. A knowledge-driven approach to text meaning processing. In *Proc. of the HLT-NAACL Workshop on Text Meaning*, 1–6. Edmonton, Alberta: ACL Press.
- Cox, M. T. 2005. Metacognition in computation: A selected history. In *Papers from the 2005 AAAI Spring Symposium*, 1–17.
- Forbus, K. D.; Gentner, D.; and Law, K. 1994. MAC/FAC: A model of similarity-based retrieval. *Cognitive Science* 19(2):141–205.
- Fox, C., and Lappin, S. 2004. An expressive first-order logic with flexible typing for natural language semantics. *Logic Journal of the Interest Group in Pure and Applied Logics* 12(2):135–168.
- Fox, C., and Lappin, S. 2005. *Foundations of Intensional Semantics*. Blackwell.
- Halpern, J. 2003. *Reasoning About Uncertainty*. MIT Press.
- Hobbs, J.; Stickel, M.; Appelt, D.; and Martin, P. 1993. Interpretation as abduction. *Artificial Intelligence* 63(12):69–142.
- Kuehne, S.; Forbus, K.; Gentner, D.; and Quinn, B. 2000. SEQL: Category learning as progressive abstraction using structure mapping. In *Proc. of the Cognitive Science Conf.*
- Lenat, D. 1995. CYC: A large-scale investment in knowledge infrastructure. *Comm. of the ACM* 38(11):33–38.
- McAllester, D. 1991. Observations on cognitive judgments. In *Proc. 9th Nat. Conf. on Artificial Intelligence (AAAI-91)*, 910–914. Cambridge, MA: MIT Press.
- McCarthy, J. 1999. Making robots conscious of their mental states. Available at <http://www-formal.stanford.edu/jmc/>.
- Moldovan, D.; Harabagiu, S.; Pasca, M.; Mihalcea, R.; Goodrum, R.; Girju, R.; and Rus, V. 2000. The structure and performance of an open-domain textual question answering system. In *Proc. of the 38th Meet. of the Assoc. for Computational Linguistics (ACL-2000)*, 563–570.
- Ngo, L., and Haddawy, P. 1996. Probabilistic logic programming and Bayesian networks. In *Algorithms, Concurrency, and Knowledge (Proc. ACSC95), Lecture Notes in Computer Science 1023*, 286–300. New York: Springer-Verlag.
- Pfeffer, A. J. 2000. *Probabilistic Reasoning for Complex Systems*. Ph.D. Dissertation, Stanford Univ., Stanford, CA.
- Poole, D. 1993. Probabilistic Horn abduction. *Artificial Intelligence* 64(1):81–129.
- Richardson, M., and Domingos, P. 2006. Markov logic networks. *Machine Learning* 62:107–136.
- Schaeffer, S.; Hwang, C. H.; de Haan, J.; and Schubert, L. K. 1993. EPILOG, the computational system for episodic logic: New user’s tutorial; Programmer’s guide; Programmer’s manual; User’s guide. Prepared for Boeing Computer Services, Seattle, WA; available at <http://www.cs.rochester.edu/research/epilog/>.
- Schubert, L. K., and Hwang, C. H. 1989. An episodic knowledge representation for narrative texts. In *Proc. of the 1st Int. Conf. on Principles of Knowledge Representation and Reasoning (KR-89)*. Extended and revised version in Tech. Rep. 345, May 1990, Dept. of Comp. Science, Univ. of Rochester, May 1990.
- Schubert, L. K., and Hwang, C. H. 2000. Episodic Logic meets Little Red Riding Hood: A comprehensive, natu-

ral representation for language understanding. In Iwanska, L., and Shapiro, S. C., eds., *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*. Menlo Park, CA, and Cambridge, MA: MIT/AAAI Press. 111–174.

Schubert, L. K., and Tong, M. 2003. Extracting and evaluating general world knowledge from the Brown corpus. In *Proc. of the HLT-NAACL Workshop on Text Meaning*. 7–13.

Schubert, L. K. 2000. The situations we talk about. In Minker, J., ed., *Logic-Based Artificial Intelligence (Kluwer Int. Series in Engineering and Computer Science, Vol. 597)*. Dordrecht: Kluwer. 407–439.

Schubert, L. K. 2002. Can we derive general world knowledge from texts? In *Proc. of the 2nd Int. Conf. on Human Language Technology Research (HLT 2002)*, 94–97.

Schubert, L. K. 2004. A new characterization of probabilities in Bayesian networks. In *Uncertainty in AI: Proc. of the 20th Conf. (UAI 2004)*, 495–503.

Schubert, L. K. 2005. Some knowledge representations and reasoning requirements for self-awareness. In *Papers from the 2005 AAAI Spring Symposium*, 106–113.

Ureel, L. C.; Forbus, K.; Riesbeck, C.; and Birnbaum, L. 2005. Question generation for learning by reading. In *Proc. of the AAAI Workshop on Textual Question Answering*.

Winston, P. H. 1975. Learning structural descriptions from examples. In Winston, P. H., ed., *Psychology of Computer Vision*. New York: McGraw-Hill. 157–209.