Information Access: Do You Mind?

John Cugini, ITL

What's the Problem?

Successful execution of many information-based tasks depends crucially on contextual knowledge. Language processing is particularly sensitive to context and I will concentrate on it in this abstract as the example *par excellence* of knowledge-dependent information access.

The only type of system currently capable of deep semantic knowledge is the human mind. This system has two notable and perhaps related aspects: 1) knowledge does not appear to be *symbolically* represented in any straightforward way - no one expects to find logical propositions directly encoded in neurons. 2) human-based knowledge appears to be intimately connected to consciousness - a phenomenon about whose nature there is simply no consensus (e.g. see Daniel Dennett vs. David Chalmers). Several development paths (not mutually exclusive) for intelligent access suggest themselves:

Muddle Through

Does successful information access really require a rich world model to provide context, or can we get by with brute force (think Deep Blue) and/or statistical algorithms? Or perhaps we can get by with many "little" scope-limited applications (e.g. airplane reservations - in which the "world model" is essentially a database system)?

It seems doubtful to me that a truly robust speech recognition system (one with performance comparable to a human transcriber) or even a text-based question-answering system could be crafted in the absence of a more general knowledge base. It's one thing to play chess using special chips and brute-force; quite another to interpret English.

Formal Representation

If extensive knowledge is required, can it be represented in a symbolic, formal (and hence manipulable) way, a la CYC (see www.opencyc.org)? Is there a "critical mass" of such common sense knowledge? Must this system be updated manually or might it achieve NLP capability sufficient to let it learn new information (including the formation of new concepts) directly from text? Just because humans apparently don't represent knowledge this way doesn't mean that it can't be done (airplanes don't need ot flap their wings to fly).

It's difficult to predict how fruitful this approach will be. It at least takes the problem of common sense knowledge seriously. But it's a big world with a lot of facts - suppose you ask a friend about her round of golf and she e-mails back: "Oh, it was pouring for a while - the clubs were practically flying out of my hands." What would it take to build a system capable of concluding that her score was probably higher than usual? Or explaining why the clubs were "flying"?

Non-symbolic approaches

Symbolic systems also have a problem of anchoring: even if the knowledge base correctly encodes the *relationships* among concepts, there must be a set of base concepts that are not further defined linguistically. Although the symbolic approach enables reasoning at the propositional level, at best it can encode conceptual knowledge only by mapping words to other words.

But humans have, for instance, a irreducibly non-linguistic understanding of what "yellow" means; human conceptual knowledge is intimately tied to sensory and emotional states (consciousness rears its ugly head), which we can label ("yellow", "pain", "sorrow") but not further explicate. And surely there are many (most?) questions the answering of which requires such non-verbal understanding.

What are the prospects for building a system that associates sensory inputs with concepts? Computer-based "neural" networks can be trained to recognize textual characters - is this simply a good trick or does it scale up to recognition of a wide variety of objects?

Conclusion

Enough smart people have been trying for a long enough time that one may safely conclude that the current limitations of intelligent access are not accidental; it will take more than yet another clever algorithm to transcend them. A significant world modeling capability is necessary to achieve major improvements over the current state of the art. What the payoff curve looks like (system performance as a function of richness of representation) is one of the major open questions.