

Panel

Next Steps in MT Research

Views from the Government, researchers, and users. This is the penultimate event of the conference, the time to think about the future. We have taken some care to compose a panel with a mix of views, backgrounds, and visions of where MT—in particular, MT research—should go from here. Should MT go after the new opportunities presented by the Internet? Or should it further mine the possibilities of statistical and other automated techniques? Should it focus instead on those difficult core issues of semantics, discourse, and pragmatics? Should it foster higher visibility to encourage common use, or should it hunker down until quality improves? Or all of these? Perspectives from research funding agencies, leading universities, and corporate research provide perspectives.

Panel members:

Lynn Carlson, U.S. Department of Defense (moderator)

Jaime Carbonell, Carnegie Mellon University

David Farwell, New Mexico State University

Pierre Isabelle, CITI

Jackie Murgida, MT contractor and system developer

John O'Hara, Globalink

Dekai Wu, Hong Kong University of Science and Technology

Observations Concerning Next Steps in MT Research

David Farwell

Computing Research Laboratory, New Mexico State University

Introduction

It appears to me that there are too many directions in which MT research might move and too many factors influencing the choice of a direction for me to prognosticate with any degree of confidence which one it will go in. Rather, I will suggest a method for picking out your own future research direction. Essentially, that is:

- to identify a basic research problem which is clearly relevant to one of two MT research goals,
- to identify the characteristics of a fundable research project given the current funding climate,
- to cast the research problem in terms of the characteristics for a fundable project.

I will predict that 'the MT problem' will neither be solved in the near future (especially since it is not an object of current interest within the MT research community) nor will it evaporate. Along the way I will make a few qualitative observations about current MT research and MT research infrastructure.

Background

The MT Problem is that translation is at bottom a cognitive process. It requires knowledge about people (speakers/authors, addressees) of both general and specific nature and about the topics being discussed of both a general and specific nature. It requires the ability to infer connections between things that are said and between why people are saying them. This in addition to assigning meanings to expressions or expressing in language meanings. The problem has certainly been known since the beginning of machine translation research in the late 1940's (best expressed by Bar-Hillel in 1960) and all research, whether wittingly or unwittingly has been directed at overcoming this problem.

This problem stands as the central barrier to progress on the two 'traditional' goals of MT research:

(1) to design and implement FAHQMT systems which are:

- efficient,
- general purpose,
- easily applicable,
- adaptable to new domains and text types,
- easy to build (and, thus, adaptable to other languages).

(2) to design and implement translation technologies (HAMT, MAHT) that support translators in performing their task more efficiently.

In pursuit of these goals, there has also been an on going tension between two broad approaches to system design and implementation. On the one hand, there is the strategy of first understanding (modeling) the process of communication (language understanding and production) under the assumption that MT system design and implementation will follow relatively transparently. On the other, there is the strategy of designing and implementing a core system first and then revising and extending that system in face of performance results (or, possibly, emerging computational techniques). The assumption here is that the overall computational model will emerge as the system's coverage and accuracy improves. Whether or not some 'deeper' model of the translation process is uncovered is not of immediate concern.

The former has been referred to as the theoretical approach, the latter as the engineering approach. In reality neither approach has been followed to the exclusion of the other (except, perhaps, in the Eurotra project).

Identifying Research Problems

With respect to the first MT research goal, FAHQMT, the focus should be on issues related to semantics and/or pragmatics. These include such issues as:

- how to represent (required) knowledge,
- how to populate (required) knowledge bases,
- how to access facts relevant to (required) inferencing,
- how to infer (required) connections between statements and their contexts,
- how to revise contexts in the face of a new statement,

- how to formulate and maintain hypotheticals or counterfactuals to assist (required) inferencing,
- how to (analogically, simulatively, abductively, deductively) reason efficiently,
- how to resolve ambiguity (syntactic, semantic),
- how to resolve reference (unique, anaphoric, cataphoric),
- how to recover ellipated information,
- how to interpret metonymy or metaphor,
- how to avoid unwanted ambiguity,
- how to avoid unsuccessful reference,
- how to avoid opaque ellipsis,
- how to avoid uninterpretable metonymy or metaphor.

The former might be cast as issues related to pragmatics, the latter as issues related to specifically grammatical (e.g., lexical and compositional semantic) processing, although that depends on the scope of these areas and the specific form suggested.

Toward the second MT research goal, translation technologies, the focus should be on issues such as:

- what are the actions (physical, cognitive) that translators perform when they translate,
- what is the physical context (setting, objects manipulated) of those actions,
- what is the physical input to the process and output from the process,
- how can existing technologies be (adapted and) applied to reduplicate those actions in that context,
- how can existing technologies be (adapted and) applied to enhance those actions in that context.

These issues (whether related to the first goal or the second goal) are independent of the methodological approaches described above. They are equally independent of the translation model used (direct, transfer, interlingual; example-based, statistical).

Identifying Characteristics Of Fundable Projects

There is one clear requirement in today's research environment on fundable projects:

project outcomes must be objectively evaluated.

This implies (although does not necessarily entail) that projects must provide quantitative results which can be evaluated objectively. The belief is that, assuming that there are no obvious inadequacies in the 'experimental' set up, success (or failure) to achieve expected results (generally results that are an improvement over prior results) indicates (1) that the hypothesized solution is headed in the right (or wrong) direction and (2) that overall progress (or lack thereof) toward some goal has taken place.

If projects must provide quantitative results, some component of the project must be directed at developing something that does something (i.e., that produces quantitative output which can be used for the objective evaluation).

The research problem, then, needs to be cast in terms of mechanisms and performance. It needs to be a tractable problem which lends itself to a quantitative definition. For many, if

not all, of the research problems listed above this recasting of the problem, both for making it computationally tractable and for providing quantifiable results, is less than obvious.

This paradigm lends itself to projects which concern statistically based attacks on large-scale, open-ended, superficial tasks or on projects which concern knowledge or inferential based attacks on small-scale, limited-scope tasks involving deeper analysis. The former are most appropriate if the overall MT application tolerates approximate, imperfect results. The latter are more appropriate if the eventual MT application requires exact, correct results. In either case, however, evaluable results can be provided.

Concluding Observations

Good research dollars are spent on projects which clearly support long term goals of FAHQMT or translation technologies that enhance the performance of translators. That is, the gains to be reaped by the overall application given success in terms of the immediate project goals must be clearly established at the outset. Otherwise, the knowledge gained is simply 'interesting'.

Bad research dollars are spent on projects whose goals are either irrelevant to the long term MT research goals or are overly trivialized specifications of goals that are relevant. Research into TL 'glossers' of SL text (or utterances) might be offered as an example of the former type. Current efforts related to word sense disambiguation and coreference resolution may serve as examples of the latter type.

Ten Years' Experience Talks about the Future!

Jackie Murgida

MT contractor and system developer

We have been asked to think about the future and discuss where MT research should go from here. For me, this is like getting an Olympic gold medal. I've been thinking about and studying MT and computer-assisted translation for over ten years, and now I'm finally being asked to tell you what I think!

I believe that there is no single approach or avenue that should be followed or even emphasized, but that all the possibilities should be explored to enable development and use of MT systems in the most beneficial way possible. Statistical techniques, example- and knowledge-based MT, corpus-based research, semantics, discourse, pragmatics, and so on—all should be explored and we should be building hybrid systems with emphasis on all these different aspects to try to find the optimal mix and balance to improve MT output.

Where I would like to see much more emphasis, because of my interest in the end users of MT, is in the area of identifying and understanding functional requirements from the user's point of view and developing improved interfaces and tools for the user to deal with the MT system itself and for the user to enhance the quality and usefulness of the raw output. This is not a new idea. A lot of good work is going on in the area of translators' workstations. I would just like to underscore the need for the MT world to keep this work up front—not as a back-burner, nice-to-have after the serious linguistic research and development is done. We need to know more about how translators and other users work with foreign-

language texts and how an MT system as a whole or modules within a system can be designed so that the user can feel comfortable with the system and get the most out of it.

In other words, regardless of how sophisticated my core MT system is linguistically, and regardless of how accurate the raw MT output is, can I sit down and postedit the output easily? Do I need a Ph.D. in computational linguistics to update the dictionary? Can I even update the dictionary? Are there options so that I can tailor the interface to my own workstyle and my own needs? For instance—and this may seem very mundane, but it's important if I have to sit here for hours using this system—do I have the choice of where the source- and target language-windows are oriented on the screen? I want that choice. And I want my choice of how unknown words are passed through the system when translating from non-latin character languages: transliterated? in the original script? if transliterated, how about having the choice of transliteration systems? We have to do more in the area of understanding how human translators translate and how non-translators are now dealing with translated texts—and with poorly-translated material or raw MT output. We need to identify the different ways individuals use existing systems (whether automated or not) and incorporate this knowledge into new systems, so that users don't have to learn a completely different way of working that's dictated by their new system.

Not only should we seek to understand different individual workstyles, we should also be looking at how people differ in their needs from an MT system according to their level of competence in a given language pair. Most of the time people in our field limit the discussion of human translators to expert humans, as if they were the only kind, suddenly appearing with no time for training and practicing, and as if there were only one level of expertise. A highly competent human translator with a native knowledge of the target language and excellent writing ability in it, along with native or near-native reading ability in the source language, will, of course, use a system differently from a novice translator, with all the variations in between. The novice might well feel more comfortable post-editing raw MT output, while the expert might prefer to translate most documents from scratch using modules of the MT system and add-on tools to facilitate research on obscure terms, find standardized spellings of proper nouns, and so on.

We have to expand the availability of automated tools, such as encyclopedias, thesauruses, concordances, monolingual and bilingual dictionaries, style manuals, gazetteers, and other reference works. These tools should make it easy to answer a question, identify text, and drop it into a translation or other document being produced with the aid of the MT software. These tools and interface options must be dictated by the existing work habits of real users. There are some features and options that are yet to be identified or at least yet to be included in current systems. For instance, does any current system allow the user to highlight a string in the target text and then be given the corresponding string in the source text, without having to hunt for it? Could the system then provide information about the source string—from the dictionary or the parse? Could the user also get information on alternative translations for the original term (that is, ones the dictionary didn't choose)? All of these, and others I haven't thought of, would be useful to some users.

Some of these points are obvious to many of you, and, again, certainly are in the realm of the mundane—not the whiz-bang cutting edge of computational linguistics, not breaking new ground in natural language understanding and generation and so on. Some are just target-language writing tools already available for use with monolingual word processors, that I'd like to see accessible in the MT posteditor. However, I believe they are important for MT systems that people can work with and are more commercially viable than some of the existing programs.