A roadmap for MT : four « keys » to handle more languages, for all kinds of tasks, while making it possible to improve quality (on demand)

Christian Boitet

GETA, CLIPS, IMAG

385, av. de la bibliothèque, BP 53

F-38041 Grenoble cedex 9, France

Christian.Boitet@imag.fr, http://www-clips.imag.fr/geta

International Conference on Universal Knowledge and Language (ICUKL2002), Goa, 25-29 November 2002

Abstract

Despite considerable investment over the past 50 years, only a small number of language pairs is covered by MT systems designed for information access, and even fewer are capable of quality translation or speech translation. To open the door toward MT of adequate quality for all languages (at least in principle), we propose four keys. On the technical side, we should (1) dramatically increase the use of learning techniques which have demonstrated their potential at the research level, and (2) use pivot architectures, the most universally usable pivot being UNL. On the organizational side, the keys are (3) the cooperative development of open source linguistic resources on the Web, and (4) the construction of systems where quality can be improved "on demand" by users, either a priori through interactive disambiguation, or a posteriori by correcting the pivot representation through any language, thereby unifying MT, computer-aided authoring, and multilingual generation.

Introduction[1]

The goal of this paper is not to enter into fine technical details, nor to support a particular MT system, because no current system, architecture, or approach is by itself sufficient to open the many locks which prevent MT from becoming as successful and useful as it should and could be. Rather, I would like to identify the most important locks, and to propose keys for opening them, let us dream, by the year 2010.

I have earlier proposed some analysis of past failures and successes of MT (Boitet, 1995), identified its most difficult scientific problems (Boitet, 1993), and risked some predictions (Boitet, 1991, 1994, 1996a, b), most of which have been confirmed. Here I would like to update these opinions in view of the explosive development of Internet and MT R&D in the last ten years. I would like to offer a synthetic view of the current state of affairs, of the "locks" preventing MT developers and vendors from doing better and doing more, and of the keys which could open these locks.

I will begin with a short assessment of current MT, distinguishing three main goals: MT aiming at rough translations of texts, MT aiming at quality translation of texts, and MT of speech. In the second section, I propose to use four (main) keys to open the identified locks: integration of numerical techniques, use of an anglo-semantic pivot (UNL), development of linguistic resources "à la Linux", and involvement of users to improve translation quality "on demand". I additionnally illustrate how these ideas can be put into practice, by using the UNL as pivot, and by cooperatively developing lexical and aligned corpus resources (Papillon project, http://www.papillon-dictionary.org/).

Short critical assessment of current MT

"Rough" MT for assimilation

Many MT systems are currently available, at low prices, for assimilation (that is, basic comprehension). The obtained translations are "rough" but often adequate for the intended purpose. However, they cannot be cost-efficiently revised to obtain quality translation. In short,

---

[1] This is a revised version of a paper presented at a workshop during MTS-VIII (2001). The perhaps futuristic or idealistic ideas presented here seem to match most of the goals and dreams of the UNL project !

- readers understand the gist of the information, or at least its topics,
- translators can use the result as a "suggestion", as they use suggestions from translation memories.

The main uses for assimilation MT now are for surfing the Web and for seeking information (military, economic, and scientific intelligence were early users, and demand continues in these areas).

However, the number of available language pairs is in fact very low compared to the needs. For example, D. Theologitis reported at LREC'2000 that EC-Systran had only 19 pairs after 24 years of developement, eight of them "almost satisfactory". Their web site[2] says they have more than 30 language combinations and 20 specialized domains, with English and French are present in most pairs. However, there are 110 language pairs in the EC alone.

In Japan (and similary in China), very few language pairs are offered besides English↔Japanese and English↔Chinese. Russian is offered for two or three pairs, and Thai only for English↔Thai.

Some Web sites claim to offer many language pairs, by translating through English. Unfortunately, the results are terrible. Try German→French and you will get a mostly incomprehensible jumble with English words in it.

The reason is very simple: the direct method employed for translation makes it impossible to "combine" (or join, or concatenate) two systems at any other level of representation for the unit of translation (UT) than the text itself. Even if the English intermediate translation were very good, many new ambiguities would appear. But that translation is very often grammatically and lexically incorrect, especially if the UT contains unknown words. In this case, the subsequent MT system can not work as intended, because it always expects mostly correct input.

The identified locks here are

- the cost of developing the first commercial version of a new language pair (at least 40 man.years according to the CEO of Softissimo),
- the direct approach, which makes it impossible to combine two systems without dramatically lowering the quality,
- the law of diminishing returns: each new language pair to be developed usually corresponds to a lesser need than the previous one, hence there are fewer users/buyers, all expecting to pay no more than the cost of already available language pairs.

Quality "raw" MT for dissemination

We find here specialized systems for (rare) niches, such as METEO (Chandioux, 1988), ENGSPAN, SPANAM (Vasconcellos & al., 1988), METAL (Slocum, 1984), LMT (McCord, IBM), CATALYST (Caterpillar-CMU), perhaps some LOGOS systems, etc. In Japan, we might mention ALT/Flash (the NTT system for Nikkei stock market flash reports) and perhaps some specialized systems, mostly EN-JP, used internally for translating technical manuals (AS-Transac at Toshiba, ATLAS-II at Fujitsu, CrossRoad at Nec, SHALT at IBM, Pensée at Oki, etc.).

In Europe, few such systems are now available, due to the relatively small market, and to the negative attitude of the EC and all governments towards funding quality MT since the completion of the Eurotra project. As our lab was not directly involved in that project, we feel free to say that this is actually quite unfair, because Eurotra was not even a precompetitive project -- no industrial firms and no users were involved. Rather, it was a research project, which produced some notable results. The project prompted interest in MT. It also answered the vocal critics of the introduction of Systran at the EC. (Systran was expensive, and remained almost

---

[2] http://www.systransoft.com/

unused until 1990, when the decision was made to abandon its use for production of professional quality translations. Instead, rough translations were distributed "as is" – 2,000 pages in 1989, and 40,000 pages in 1990, according to L. Rolling at MMT'90.)

References to (quality MT) systems for dissemination are very rare. They are indeed very good and very useful (30 Mwords/year for METEO, 75% EN-FR and 25% FR-EN, with 1 mn revision per page, 0.15 cents/word for final output), because they are quite specialized.

Let us remark here that it is extremely difficult to prepare comparative benchmarks for such systems, because, like expert systems, they are very good on their domain, and fail miserably on other tasks. The best way to measure them is through some combined assessment of the buying, maintenance, and evolution prices, and through consideration of the human time needed to obtain a professional result.

I have proposed elsewhere the formula "Quality x Coverage = Constant", with the constant depending on the particular system. To augment the translation quality, then, there are three combinable ways:

- specialize to a sublanguage (domain, grammar; this is known as the "suboptimization approach" (Lehrberger & al., 1988),

- involve the user on the source side (using controlled language as in KANT-CATALYST (Nyberg & al., 1992), or interactive disambiguation (Boitet & al., 1994, Wehrli, 1992),

- improve the overall approach, e.g. by introducing more abstract linguistic levels (functional, relational, logical, etc.), more "felicitous" data structures (decorated trees, typed feature structures, charts, hypergraphs, etc.), non-determinism, scoring, etc.

As the needs are real, in particular for technical 1—>N translations, there is still some

activity in the quality translation field, but far less than was hoped ten years ago.

Technically, these systems almost always have a separate analysis component, producing a syntactic or syntactico-semantic descriptor of the source UT (usually an annotated or decorated tree). Almost all use some flavor of the transfer approach (even systems like ATLAS-II by Fujitsu or PIVOT-CROSSROAD by NEC).

In most cases, there is no syntactico-semantic descriptor of the target UT, transfer and generation being merged into a single phase using recursive descent of the analysis tree. Hence, changing the source language implies redoing all the work. (See the difficulties experienced by Siemens with METAL in the 90's, which contributed to the company's exit from the scene.)

The identified locks here are

- the cost of developing the first commercial version of a new language pair: at least 100 man-years according to H. Sakaki, the main author of KATE at KDD, and to Pr. Nagao, director of the MU project, and perhaps 300 man years with large dictionaries, as H. Uchida estimated for ATLAS-II at Fujitsu,

- the impossibility of factorizing generation processes, when the situation changes from 1→N to m→N,

- the fact that, although systems with full transfer structure could be used to produce "all language pairs" by combining systems at the levels of the structural descriptors, there seem to be no industrial situations at the moment calling for high quality for many or all language pairs.

Speech translation

Current commercially available technology makes SST already possible and usable for "chat MT". Such systems are usually built by combining speech recognition (SR), text MT, and speech synthesis. NEC has demonstrated a system for JP↔EN at Telecom'99, and probably markets it. Linguatec, a

subsidiary of IBM, demonstrated Talk&Translate at COLING2000 (which uses Via Voice and LMT for EN—GE—FR). The quality is of course not very high in all components, but this drawback is compensated by the broad coverage, by some feedback (e.g., editable output of SR and written reverse translation), and by the fact that users are intelligent humans wanting to communicate.

At the research level, the aim is to obtain higher quality while allowing more "spontaneous" speech, in task-oriented situations. The large German VerbMobil project (1992-2000) has shown the feasibility of reaching these goals, and has compared many alternative methods in the same task setting (Wahlster, 2000). The goals can also be reached in a multilingual setting, as demonstrated by the CSTAR consortium in intercontinental public demos with large media coverage in July 1999. Participants used a kind of "semantico-pragmatic" pivot designed to represent the utterances of participants in a limited set of situations (e.g. exchanging tourism information, booking hotels or tickets for events and transports, etc.).

The higher quality is necessary because at least one participant (the "agent") is a professional who must work fast. S/he may adapt to the system, but still can not afford to repeat each utterance two or three times until the system understands it correctly. The higher spontaneity is necessary because at least one participant (the "client") is supposed to be a naive and occasional user of the system.

The identified locks here are

- the great difficulty of developing an adequate pivot (the IF or Interface Format of CSTAR),
- the cost of building the necessary lexical resources, as for MT of texts,
- the difficulty of handking the context (pragmatic, discursive, linguistic, lexical), in particular to compute correctly the speech acts and the referents of anaphoras and elisions.

Fundamental research is still badly needed to improve overall quality and enlarge usability, in particular on

- context processing: how to transmit and use possible "centers" (identifiable entities) usable
  - in analysis, for anaphora or elision
  - in generation, for controlling lexical selection and producing ellipses and elisions to improve naturalness and coherence
- prosody processing: how to generate prosodic marks (to be used by the TTS components) from pragmatic, semantic, and syntactic features
- integration among heterogeneous components (SR, MT):
  - richer interface data structures (such as tree lattices),
  - use of common primary linguistic resources (lexical and grammatical data bases),
  - system architecture (pipe-line, agents, blackboard, whiteboard).

Current research centers on almost fully automatic systems, leading to extremely specific, task-dependent systems. While they can be useful, we should not repeat the errors made in MT during the 70's. It is possible to develop computerized aids

- for interpreters (to help several conversations, which may in part be conducted directly in some common language, or may be conducted indirectly through some imperfect SST system)
- for "active listeners" wanting to better understand speech in a foreign language (conversation, radio, TV, etc.).

This topic is not purely technical and leads to quite a few interesting scientific questions.

Four "keys"

To open the door toward MT of adequate quality for all languages (at least in principle), we propose to use four keys.

On the technical side, we should dramatically increase the use of learning techniques which have demonstrated their potential at the research level, and should use pivot architectures, the most universally usable pivot being UNL.

On the organizational side, the keys are the cooperative development of open source linguistic resources on the Web, and the construction of systems where quality can be improved "on demand" by users, either a priori through interactive disambiguation, or a posteriori by correcting the pivot representation through any language, thereby unifying MT, computer-aided authoring, and multilingual generation.

Learning techniques

Whatever is said in conferences, the reality of current MT is that all commercial systems are rule-based, and that their linguistic resources, lexical, syntactical and semantic, are by and large hand-crafted. Of course, developers try to start from existing computerized resources, but what they obtain is far closer to unprocessed ore than to a finished product.

This is still very true of dictionaries, but there are many ongoing efforts to alleviate the problem through automatic lexical data collection (terminology extraction) and the use of exchange formats.

When it comes to the syntactic knowledge, even the patterns of EBMT (or TDMT) systems are manually derived from the collected examples. However, research or prototype systems integrating automatic learning seem quite promising. We aren't speaking here of "language models" used in SR and largely based on trigram and bigram frequencies, or of HMMs, because they are in no way powerful enough to produce structural descriptors of entire utterances, not to speak of paragraphs or entire texts.

The two directions which look the most promising are :

- the methods for learning the transitions of parsers directly producing "semantic trees" in task-oriented situations such as the querying of data-bases (Roucos, IBM, Via Voice group),

- the recent advances in research on translation memories, showing how to abstract patterns with variables from monolingual or bilingual examples.

Using non-textual pivots

It is true that the use of any sort of "pivot" with autonomous lexical symbols leads to double lexical translation. Further, precision of translation may be lost if the pivot representation of an utterance is incomplete relative to some (interlingual) features underspecified in some languages and necessary in others, such as number, gender, aspect, or modality. These problems inevitably yield a decrease in the "asymptotic quality" of automatic translations, if the criterion for quality is that of professional technical translation, namely the highest possible parallelism in form and content.

However, there are many situations -- mostly in human communication with humans or with machines, but also in technical domains — where a less stringent criterion is adequate. In these situations, paraphrasing instead of translating is acceptable, provided the informational content and if possible the communicative aspects of the source text are rendered quite exactly.

As a matter of fact, CATALYST, used by Caterpillar for the translation of technical documents, is based on very abstract representations linked with the ontology of the domain, and many of its translation examples would not be accepted by translation teachers, but are very good for the purpose at hand. The program thus demonstrates a way of combining MT and multilingual generation from abstract representations.

What are, then, the most promising kinds of pivots for the future?

Semantico-pragmatic pivots

Examples are the interlingua of CATALYST, and the interface formats used

by the Verbmobil and CSTAR speech translation projects. However, such pivots cannot really open the doors to a dramatic increase in the number of situations and languages which can be handled by future MT systems, because

- they are very costly to build, and costly to adapt to similar tasks,
- it is impossible to extend them gracefully to handle the "full" language.

This last point has been being rediscovered the hard way by the Nespole! project. The attempt to extend the CSTAR-II IF to colloquial speech has led to large but partly incomprehensible and provably inconsistent and ambiguous specifications, which still can not express the variety of naturally produced utterances. This is an absolutely fundamental point : trying to build this sort of axiomatization of the "full" language is doomed to failure, much as trying to build a complete axiomatization of all arithmetic truths (proved to be impossible by Gödel).

Abstract linguistic descriptors

For quality translation, using detailed abstract linguistic descriptors of a particular language as a "pivot" offers distinct advantages where language coverage, precision, and building costs are concerned.

- Choosing a language having many reusable resources, such as English, Chinese, Russian or French, and "central" according to the set of languages aimed at, makes it possible for all developers to understand the pivot structures.
- Use of syntactico-semantico-logical structures such as B. Vauquois' "m-structures" (multilevel structures) has been demonstrated to give the best degree of precision since the early days of MT.
- Lexical resources can be derived from available and familiar sources such as monolingual and bilingual dictionaries, whereas invented lexical symbols of other kinds of non-textual pivots offer much more room for variations in interpretation

by developers. Costs diminish and quality increases.

Anglo-semantic pivot: UNL

In my view, the most promising pivot for use as a "key" to give many languages access to the realm of MT is UNL[3], because it is in essence an "anglo-semantic pivot". The fact that all its symbols are built from English words opens UNL to all languages, because all developers of NLP systems in the world have at least a working knowledge of English. In a way, English is used internally to get rid of English — or, to put it in a more positive way, to make all languages equal and overcome the "language divide".

Another factor is that UNL is structurally very simple, whereas linguistic descriptors such as m-structures necessarily reflect the syntax of a natural language, which is always extremely rich and complex. By contrast, semantic graphs or even hypergraphs are very easy to grasp. Thus, to extend "rough MT" to "all languages", developers of assimilation-oriented MT systems could use UNL, without too much investment in the UNL technology itself, to get a "squared" result (in terms of the numbers of languages) with only a "linear" investment.

The UNL project, started by the UNU in December 1996, and opened to the public in November 1999, is the only current project offering such a viable framework. Let us say only a few more words about it here, while referring to the literature for more details (Boitet, 1999, Sérasset & al., 1999, 2000). Sixteen countries are participating. The project has always been presented as a project for multilingual communication and information retrieval over the Web, insisting that it is not YAMTP (Yet Another MT Project).

In accordance with that idea, and with the situations envisaged, enconversion (generation of a UNL graph, given a source text) is presented as not necessarily automatic, or even semi-automatic. In fact, experiments

---

[3] http://www.unl.ias.unu.edu/ or http://www.undl.org/

with students working on everyday and technical texts have shown that producing UNL graphs directly with a bare bone editor (such as BBEdit) is quite feasible. Of course, computer aids such as UNL-oriented menu-driven and/or graphical editors have been developed, and give better productivity. After a delay due to funding problems, work is also beginning on the construction of fully automatic or semi-automatic analyzers (including preediting and/or interactive disambiguation) for several languages.

About twelve automatic "deconverters" (programs which produce text, given UNL graphs) have been developed, with varying coverage (30000 to 100000 dictionary entries, very small or large grammars). Six to eight of them are accessible as Web servers.

UNL (Universal Networking Language) is

- a language of "anglo-semantic hypergraphs" used to represent the linguistic content of documents in a language independent way, or, more precisely, to represent any utterance of any language by the abstract sructure of an English utterance thought to express essentially the same meaning, though perhaps lacking some of the information which would be necessary to produce a precisely equivalent English surface expression;

- a computer file format embedded in html[4] to represent multilingual parallel documents aligned at the level of utterances as one file, each utterance having representations in the UNL language and in several natural languages.

- on international cooperative project.

UNL could thus be called "the html of linguistic content". The "linguistic content" of an utterance or a fragment is represented as a graph or hypergraph, where each arc is labelled by a semantic relation (roughly, a deep case) and each node is labeled by a "universal word" (UW) or recursively contains a

UNL graph, each graph or subgraph having an "entry node". A UW is a lexical symbol denoting a (set of) acceptions (meanings, or word senses), and is constructed by borrowing from English because all developers know English: it is an English term followed by formal restrictions enclosed in parentheses, for example "chair(icl>thing)" or "chair(icl>do, agt>human, obj>entity)".

In the UNL file format, which has associated tools linked with Internet Explorer, a document is one file, and each document is multilingual. There are special tags to delimit sentences (or fragments), and, at the next level, tags to delimit the original version, the corresponding UNL graph, and the versions in all languages (if any). This format thus provides a solution for the encoding of the multilingual content.

Cooperative development of open source linguistic resources on the Web

Due to the law of diminishing returns, no firms and no institutions are likely to invest heavily in linguistic resources for the majority of languages not yet covered by MT, and by NLP tools in general. We propose to develop these resources through a cooperative development "à la Linux", aiming at the production of shareware resources. The basic idea is that there are competent internauts "out there" who would like to help if they have some reasonable incentive, such as the free access to resources, or, even better, to tools built on these resources. Let us illustrate this plan through examination of the ongoing Papillon project.

This project was originally motivated by the lack of large French-Japanese dictionaries usable by French speakers not knowing Japanese and hence unable to pronounce and understand the Japanese parts of existing dictionaries. It was also felt that the lack of bilingual FR-JP resources is an obstacle to the development of linguistic software applications involving French and Japanese. Because of this lack, applications that have been created thus far for French and Japa-

---

[4] an equivalent xml form is used at GETA.

nese have only a limited scope, whereas good English-Japanese programs are available. And yet, Japan is certainly very interested in the French language. Conversely, a growing number of French individuals are investing considerable energy to learn Japanese. There is thus a vacuum to be filled.

It was realized that the Thai student community in Japan had encountered exactly the same problem and had recently produced a Web environment (the SAIKAM projet, a joint undertaking of NII, Tokyo, and NECTEC, Bangkok) to build a Japanese-Thai on-line dictionary, modifiable at any time by its users.

A similar project for English and Japanese has been active for about a decade. This JMDIct project has allowed the effective construction of a free Japanese-English dictionary, available through an Internet server (Pr. Jim Breen, Monash University, Australia, http://www.csse.monash.edu.au/~jwb/edict.html). The current dictionary comprises 70,000 entries of common vocabulary, a specific kanji dictionary, and around 20 specialized dictionaries (biology, law, etc).

Thus the Papillon project, a similar grass-roots bridge between French and Japanese, has been proposed. The project uses a "pivot" architecture. Monolingual dictionaries follow the "DiCo" format, simplified from the DEC format but still very rich (Polguère & Mel'tchuk). The units of DiCos are "lexies", or senses of "vocables". The DiCos are linterlinked by "axies" grouped in a central structure. An axie has a list of lexies for each natural language, links to the symbols used by other represenation systems (WordNet, UNL, Lexiguide, EDR, etc., the list of systems being open), and semantic links corresponding in some way to the lexico-semantic functions linking lexies in a DiCo.

The project's server is open for consultation and will open for contributions in 2002-03. To use it or contribute to it, one has to open an account. The main principle is that the data are open source, in the Linux spirit, and that each contribution is credited to its author. Anyone can become a contributor at any time by entering some information relative to a vocable, a lexie, or links, or by providing remarks, or even files in some known format (such as DEI). The user's contributions will be put in his or her reserved space, in the form of XSLT transformations. Also, groups will be defined by the users, or communities of users, which belong to them, with appropriate access rights. To ensure quality, only the predefined "central group" is allowed to validate and possibly correct the contributions before putting them in the common base.

For several years, we have also been trying to set up a similar project (Montaigne) for mutualizing translation memories and tools for translators. The idea is that occasional translators from various communities (such as people working on MT, on parallel computing, on AIDS research, etc.), can share a large translation memory and translation aids, all residing on a server, provided that they are willing to share the result of their work (aligned sentences, fragments, and terms). Past obstacles have been the proprietary nature of available tools (Eurolang Optimizer in the beginning, then TM-2, Trados WB, Transit, and XMS) and the difficulty of programming an editor of bitexts, running on a server and using only available browsers as clients on any platform. But these problems have been overcome (see http://www.yakushite.net/ by OKI and http://www.laosoftware.com/ by V. Berment).

Construction of systems where quality can be improved "on demand" by users

The capacity to improve quality only when the user deems necessary is an essential feature of multilingual authoring and generation systems such as MULTIMETEO (Coch & al., 2001). Until now, no MT system has offered this possibility.

An important advantage of UNL is that the quality of the UNL graph corresponding to an utterance can be improved by modifying the graph in relatively simple ways, both a priori or a posteriori.

A priori means before deconversion (construction of a target text, given a UNL graph). Of course, interactive disambiguation is possible, but there is nothing special about UNL in this respect. What is possible with UNL and not with other types of representations considered for MT is that

- UNL graphs are quite compact and understandable compared with other standard linguistic structures (especially those of HPSG and the like!),

- it is possible to present the graphs in the source language, as demonstrated by UNL-Spain,

- it is possible to design user-friendly interfaces for helping users to build UNL graphs from scratch, or to modify existing graphs.

A posteriori means after deconversion (generation of target text) has occurred. Suppose we get an output text in Spanish, with almost all of the articles wrong, with number errors (singular/plural) here and there, etc. Such errors could easily arise if the text had been deconverted from a UNL graph produced from Chinese, Russian, Japanese, or Thai, in which information concerning determinateness, number, etc. is not explicit. The solution, then, involves establishment of a correspondence between an input text (here, in Spanish) and the associated UNL graph. For this purpose, one can use a Spanish-UNL dictionary, or, if one is not accessible, a Spanish-English dictionary, because the UWs (UNL lexical symbols) are built from English "heads" and semantic restrictions. It should then be possible to build interfaces allowing users to modify the UNL graph without even seeing it, "through the natural language". Techniques for modifying an abstract interlingual structure via NL expressions already exist and are used every day

(see the MULTIMETEO system (Coch & al., 2001), covering six languages).

After the UNL graph is modified, the Spanish user might even have the satisfaction of contributing to the quality not only of the Spanish utterance, but of the corresponding utterances in French, Italian, German, etc. (if s/he knows these languages).

Hence, the linguistic quality of a multilingual document (in UNL format) could be increased "on demand", piece by piece, by people reading it in different languages. This is the subject of on angoing Ph. thesis, and first versions of such a « coedition » environment have been presented at LREC and COLING in 2002.

Conclusion

To open the door to MT of adequate quality to all languages (at least in principle), we have proposed to use four keys.
On the technical side, we should

- dramatically increase the use of learning techniques, be they symbolic, numerical, or mixed, which have demonstrated their potential at the research level, and

- use pivot architectures, the most universally usable pivot being UNL.

On the organizational side, the keys are

- the cooperative development of open source linguistic resources on the Web, and

- the construction of systems where quality can be improved "on demand" by users, either a priori through interactive disambiguation, or a posteriori by correcting the pivot representation (UNL or other) through any language.

An underlying theme is to merge MT, computer-aided authoring, and multilingual generation, through user-friendly « coedition » environments.

On the practical side, we should also seek keys to unlock private investment and/or public funding !

References

Boitet (1990) Multilingual Machine Translation does not have to be saved by Interlingua. Proc. MMT'90, 5-6 Nov. 1990, 2 p.

Boitet (1991) Quelle automatisation de la traduction peut-on souhaiter et réalitations de travail individuelles ? Proc. Colloque "L'environnement traductionnel. La station de travail du traducteur de l'an 2001", avril 1991, Presses de l'Université du Québec, pp. 3—19.

Boitet (1993) Crucial Open Problems in Machine Translation and Interpretation. Proc. Symposium on Natural Language Processing in Thailand, 17-21 March 1993, Chulalongkorn University, pp. 1—29.

Boitet (1994) Dialogue-Based MT and self-explaining documents as an alternative to MAHT and MT of controlled languages. Proc. Machine Translation 10 Years On, 11-14 Nov. 1994, Cranfield University Press, pp. 22.1—9.

Boitet (1995) Factors for success (and failure) in Machine Translation — some lessons of the first 50 years of R&D. Proc. MTS-V (Fifth Machine Translation Summit), 11—13 July 1995, CEE, 17 p.

Boitet (1996a) (Human-Aided) Machine Translation: a better future? In "Survey of the State of the Art of Human Language Technology", Cole (Editor-in-Chief) & al., ed., Giardini, Pisa, pp. 251—256. (also available since 1996 at http://www.cse.ogi.edu/CSLU/HLTsurvey/)

Boitet (1996b) Machine-Aided Human Translation. In "Survey of the State of the Art of Human Language Technology", Cole (Editor-in-Chief) & al., ed., Giardini, Pisa, pp. 257-260. (also available at http://www.cse.ogi.edu/CSLU/HLTsurvey/)

Boitet (1999) A research perspective on how to democratize machine translation and translation aids aiming at high quality final output. Proc. MT Summit VII, Singapore, 13-17 September 1999, Asia Pacific Ass. for MT, pp. 125—133.

Boitet & Blanchon (1994) Multilingual Dialogue-Based MT for Monolingual Authors: the LIDIA Project and a First Mockup. Machine Translation, vol. **9**/2, pp. 99—132.

Boitet & al. (2000a) A way to integrate context processing in the MT component of spoken, task-oriented translation systems. Proc. MSC-2000, Kyoto, 11-13/10/2000, ATR, 5 p.

Boitet & Guilbaud (2000b) Analysis into a formal task-oriented pivot without clear semantics is best handled as "usual" translation. Proc. ICSLP-2000, Beijing, 16-19/10/2000, CMFP, vol. 4/4, pp. 436—439.

Chandioux (1988) 10 ans de METEO (MD). In "Traduction Assistée par Ordinateur. Actes du séminaire international sur la TAO et dossiers complémentaires", Abbou, ed., Observatoire Francophone des Industries de la Langue (OFIL), Paris, mars 1988, pp. 169—173.

Coch & Chevreau (2001) Interactive Multilingual Generation. Proc. CICLing-2001 (Computational Linguistics and Intelligent Text Processing), Mexico, Springer, pp. 239-250.

Lehrberger & Bourbeau (1988) Machine Translation. Linguistic characteristics of MT systems and general methodology of evaluation. John Benjamins, 240 p.

Levin & al. (1998) An Interlingua Based on Domaine Actions for Machine Translation of Task-Oriented Dialogues. Proc. ICSLP'98, 30th November - 4th December 1998, Sydney, Australia, vol. 4/7, pp. 1155-1158.

Nirenburg (1989) Knowledge-based Machine Translation. (Machine Translation), vol. /4, pp. 5-24.

Nyberg & Mitamura (1992) The KANT system: Fast, Accurate, High-Quality Translation in Practical Domains. Proc.

COLING-92, 23-28 July 92, Boitet, ed., ACL, vol. 3/4, pp. 1069—1073.

Sérasset & Boitet (1999) UNL-French de-conversion as transfer & generation from an interlingua with possible quality enhancement through offline human interaction. Proc. MT Summit VII, Singapore, 13-17 September 1999, Asia Pacific Ass. for MT, pp. 220—228.

Sérasset & Boitet (2000) On UNL as the future "html of the linguistic content" & the reuse of existing NLP components in UNL-related applications with the example of a UNL-French deconverter. Proc. COLING-2000, Saarbrücken, 31/7—3/8/2000, ACL & Morgan Kaufmann, vol. 2/2, pp. 768—774.

Slocum (1984) METAL: the LRC Machine Translation system. In "Machine Translation today: the state of the art (Proc. third Lugano Tutorial, 2–7 April 1984)", King, ed., Edinburgh University Press (1987).

Sugaya & al. (1999) End-to-End Evaluation in ATR-MATRIX Speech Translation System between English and Japanese. Proc. EUROSPEECH'99, September 5-9, 1999, Budapest, Hungary, vol. 6/6, pp. 2431-2434.

Sumita & al. (1999) Solutions to Problems Problems Inherent in Spoken-language Translation: The ATR-MATRIX Approach. Proc. MT-Summit VII, Singapore, September 13-15, 1999, vol. 1/1.

Vasconcellos & León (1988) SPANAM and ENGSPAN : Machine Translation at the Pan American Health Organization. In "Machine Translation systems", Slocum, ed., Cambridge Univ. Press, pp. 187—236.

Wahlster (2000) Verbmobil: Foundations of Speech-to-Speech Translation. Artificial Intelligence, Springer, Berlin, Heidelberg, New-York.

Wehrli (1992) The IPS System. Proc. COLING-92, 23-28 July 1992, Boitet, ed., vol. 3/4, pp. 870-874.

-o-o-o-o-o-o-o-o-o-o-